AD________________

Award Number: DAMD–17–03–2-0001

TITLE: Advanced Technologies in Safe and Efficient Operating Rooms

PRINCIPAL INVESTIGATOR: Adrian E. Park, MD

CONTRACTING ORGANIZATION: University of Maryland Medical Center
Baltimore, MD 21201

REPORT DATE: February 2007

TYPE OF REPORT: Annual

PREPARED FOR: U.S. Army Medical Research and Materiel Command
Fort Detrick, Maryland 21702-5012

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*
01/02/07

2. REPORT TYPE
Annual

3. DATES COVERED *(From - To)*1
Feb 2005 – 28 Feb 2007

4. TITLE AND SUBTITLE

Advanced Technologies in Safe and Efficient Operating Rooms

5a. CONTRACT NUMBER

5b. GRANT NUMBER
DAMD17-03-2-0001

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Adrian E. Park, MD

E-Mail: apark@smail.umaryland.edu

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

University of Maryland Medical Center
Baltimore, MD 21201

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
U.S. Army Medical Research and Materiel Command
Fort Detrick, Maryland 21702-5012

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT
Approved for Public Release; Distribution Unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT: The three major research targets of this study are (a) OR informatics (b) simulation research, and (c) smart image. The purpose of the OR informatics program is to develop, test, and deploy technologies to collect real-time data about key tasks and process elements in clinical operating rooms. The objective of Simulation research is to create a system where a user can interact with a virtual human model in cognitive simulation and have the virtual human respond appropriately to user queries and interventions in clinical situations, with a focus on cognitive decision making and judgment. The objective of smart image is to use real-time 3D ultrasonography and 40-slice high-frame-rate computed tomography (CT) for intraoperative imaging to volume rendered anatomy from the perspective of the endoscope. The overall project reported here was modified to add and additional task related to the DARPA-sponsored, TraumaPod project. The period of this contract was extended to February 28, 2008. This extension was due to the delay in obtaining IRB approval and other administrative matters. Based on the extended period of performance, this project is on time, on schedule, and within performance parameters.

15. SUBJECT TERMS
No subject terms provided

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			USAMRMC
U	U	U	UU	54	**19b. TELEPHONE NUMBER** *(include area code)*

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

Table of Contents

Page

INTRODUCTION 1

BODY 1

KEY RESEARCH ACCOMPLISHMENTS...10

REPORTABLE OUTCOMES 14

CONCLUSION...16

REFERENCES 18

APPENDICES..20

Introduction

The three major research targets of this study are (a) OR informatics (b) simulation research, and (c) smart image. The purpose of the OR informatics program is to develop, test, and deploy technologies to collect real-time data about key tasks and process elements in clinical operating rooms. The objective of Simulation research is to create a system where a user can interact with a virtual human model in cognitive simulation and have the virtual human respond appropriately to user queries and interventions in clinical situations, with a focus on cognitive decision making and judgment. The objective of smart image is to use real-time 3D ultrasonography and 40-slice high-frame-rate computed tomography (CT) for intraoperative imaging to volume rendered anatomy from the perspective of the endoscope. The overall project reported here was modified to add and additional task related to the DARPA-sponsored, TraumaPod project. The period of this contract was extended to February 28, 2008. This extension was due to the delay in obtaining IRB approval and other administrative matters. Based on the extended period of performance, this project is on time, on schedule, and within performance parameters.

Body

The content of this annual report contains information pertinent to continued activities in relation to the DAMD–17–03–2-0001 (Modified) “Advanced video technology for safe and efficient surgical operating rooms” project (“ORF-Y3”). This contract consists of a scope of work that fits seamlessly onto the current continuing research and activity in the contract W81XWH-06-2 “Advanced technologies in safe and efficient operating rooms” work (“ORF-Y4”). The ORF-Y3 activities have consisted of the three areas of research: (a) OR informatics (b) simulation research, and (c) smart image, plus (d) the TraumaPod project, and (e) video-assisted coordination tools. Objectives listed in this annual report are aligned with current (Y4) objectives. Where specific objectives are indicated, they reference objectives in the ORF-Y4 contract.

A. OR Informatics

Informatics subgroup 1. Workflow and Operations Research for Quality (WORQ)

In accord with the proposed timeline, preparatory analysis and systems evaluation is nearing completion. This project has been focusing on the technology assessment of evaluating active RFID technology. We have selected a technology stack based upon a “wireless mesh” network, and we have initiated the procurement process. Using this system will avoid the extensive costs and delays associated in laying the network infrastructure in the Operating Rooms. The ease of deployment became a critical requirement due to the difficulty and cost in taking an operating room out of commission. A wireless mesh network relays signals off of each station to a base receiver station connected to the network. This allows us to deploy and realign an RFID-sensor network quickly, as stations only require an AC power source but no network access. We also feel that this technology is particularly apt for a mobile military application, in which there is minimal infrastructure needed for deployment. The solution we selected had two additional advantages: First, the system uses the Zigbee protocol which does not collide with existing 802.11 bandwidths which are used for other purposes. Second, we were able to establish an engineering relationship with the vendor, and received an associated developmental pricing structure. The vendor has an excellent messaging architecture and is allowing us access to the Advanced Programming Interface (API). The system will use event triggering to send us messages to our context awareness system. The vendor is very excited about extending the functionality of their system and finding new uses for RFID in the OR. This partnership allows for a straightforward commercialization of any products created by our research team.

Additionally, we have recruited two graduate students from the computer science department of UMBC. They have been provided space at the Simulations Center, and are currently carrying out preliminary work.

Informatics subgroup 2. Operating Room Glitch Analysis (OGA)

The OGA project, focusing on institutional learning, is looking at the workflow around performance indicators in the perioperative environment and building a graphical dashboard to allow data mining and trend analysis of operating indicators.

This project started in November with a kickoff meeting on November 17, 2006. At that meeting, we developed several clinical scenarios with key stakeholders and subject matter experts. We developed a relationship with the clinical IT team to get access to the clinical scheduling and tracking system in the department. We established ties with the office of clinical effectiveness for the hospital and worked with their data analysts and nurse coordinators to understand the failure modes associated with clinical cases start time delays.

We have assessed the data source and data collection methodology for the performance data. We will be using standard ETL (extraction, transform, and loading) of the performance data. The database system and server have been setup has been established in compliance with all HIPAA and security practices associated with PHI. The lead web developer and database designer tasks were given to Max Warnock.

We recruited a logistics professor from college park, Dr. Jeffrey Herrmann, with extensive experience in lean process improvement. Dr. Herrmann has begun his work with our group, and is building a root cause analysis of the types of waste associated with delays.

Dr. Yan Xiao conducted an exhaustive literature review of peer reviewed publications on operating room activity analysis, and disseminated resources via a web-based repository for the team to review.

Informatics subgroup 3. Context Aware Surgical Training (CAST)

Objective 1.1.

In 2006, student Sheetal Agarwal, collaborated with Dr. Geoffrey Kuzmich to identify a scenario in the Trauma Surgery Training area. She created a system to automatically populate the MER (Medical Encounter Record). It was designed to be able to infer high level events such as hypovolemia from low level events in the ORF and from streaming vital sign data coming from medical monitoring systems. The streaming vital sign data came from a Perl script that read a data file that came from Stan, the HPS METI patient simulator. These clinically significant events could then be used to tag the (training) video of a procedure.

Objectives 1.2 and 1.3.

The fall of 2006 began with the preparation of a feasibility demonstration that incorporated the work of previous students. The demonstration consisted of using RFID tags to identify the arrival of a surgeon, a patient and an electronic consent form to the operating room. The distributed application would announce the arrival of each of the aforementioned, and for the surgeon and patient would display information that was stored in a database. As each agent was detected, the application would announce what or who was still missing in order to proceed and start the procedure. When the application detected all were present, it would turn off the lights and start the MER system.

The demonstration was intended to show the potential of pervasive computing in the Operating Room of the Future (ORF). In addition to the previously mentioned functionality, such an environment would also be able to detect when a patient is delivered to an incorrect operating room, when an incorrect blood type for a patient is in the ORF, or when the wrong set of tools for a given procedure have entered the ORF. Simple voice commands would be used to log information about the procedure.

Over the course of the semester, the MER has been developed into a secure electronic medical record of the events in the operating room with the purpose of alleviating the work load of medical staff so they can better focus on the task at hand thereby reducing errors and increasing efficiency.

This demonstration was presented at the UMBC campus. We then brought it to UMMC on October 4th and configured it to run in the University of Maryland Simulation and Technology Center for a public event attended by the media. The same demonstration was also presented at the official opening of the Center on December 6th.

Graduate assistants also have been observing the workflow of procedures in the operating room and talking to staff to better understand the dilemmas facing medical staff in the operating room. We anticipate that the students will begin meeting with Drs. Kavic and Turner in January to work on the requirements for the project. We will investigate if we can acquire data from the new laparoscopic training system to incorporate in the CAST demonstration system.

Informatics subgroup 4. Operating Room Clutter (ORC)

ORC Objective 1: Establish a quantitative, valid measure of workflow in the patient-surgeon interfaces.
A set of preliminary measures was developed and piloted tested on archival data from 3 minimal invasive procedures.

ORC Objective 2. Identify ergonomic problems as consequences of workplace designs (such as arrangement or management of cables and catheters)
An initial set of ergonomic problems associated with cords and cables used during minimal invasive procedures was established and incorporated into the coding scheme for future analysis of procedures.

ORC Objective 3. Collect a set of illustrative still images and short video clips to demonstrate key barriers of optimal workflow that have direct safety and efficiency concerns
Using archival video records as data, case analysis techniques were piloted to determine the effectiveness of coding schemes. Using the pilot methods, seven cases of minimal invasive procedures were tentatively analyzed.

ORC Objective 4. Develop conceptual future workplace layouts that optimize patient-surgeon interfaces
No progress on this objective. We expect to make progress on this objective following completion of data collection of phase 1.

Additional Objective: Literature collection and review:
Literature or surgical ergonomics, operating room layout, and analysis of operations-related issues in OR management has been collected and cataloged.

Informatics subgroup 5. Improving Perioperative Communications (IPC)

The progress in the IPC section has been impeded by the departure of a key personnel resource from the University of Maryland Medical Center, who provided both access to and information regarding operational data. While the IPC activities have continued and will continue with alternate personnel, the personnel change has slowed progress.

Objective 1. Pre-planning
Pre-planning has been completed, with the establishment of high-level process mapping of communication, and preliminary needs assessment to determine key players and stakeholders.

Objective 2. Determine necessary information to be conveyed and appropriate individuals
Initial analysis of information requirements has been undertaken. In conjunction with Objective 3 (scripting for the phone line), the sufficiency of this analysis is being examined.

Objective 3. Develop scripts and phone line
Initial revision of an improved scripting has been carried out for the scheduling process phone line.

Objective 4. Develop technological components of accessing phone logs and usage information.
Initial examination of available operational data was carried out. Analysis of data revealed a number of deficiencies in the specificity of the currently-available data. Specifically, delay codes available in the current information system do not provide sufficient detail in their current format. Delay annotations, written in free text, are sometimes available, but these data cannot be analyzed using automated reporting systems standard data analysis techniques. Work being carried out in the Informatics Subgroups 1 and 2 may provide additional data which will facilitate the acquisition of relevant data.

Administrative note: The Informatics sub-projects of WORQ, IGA, and Improving Perioperative Communications projects were all delayed significantly in their progress by the departure of a key resource from the University of Maryland Medical Center (Patricia Smale). This individual served as a clinical nurse specialist with oversight of OR activities, access to operational data, and influence on the logistics of implementation of the research infrastructure. No one has yet been hired to fill the position. The sub-groups are compensating for this loss by working with other individuals in the interim to serve the same function for our research.

B. Simulation

Objective 1. Coverage of diseases for simulation and treatment will be extended to include three to five additional diseases (depending on the depth of description, to be determined by the team of MDs working with the technologists).
We have added three new diseases to the simulation capabilities of the Maryland Virtual Patient (MVP): Zenker's diverticulum, scleroderma-esophagus, and LERD (the latter is partially implemented).

Objective 2. Simulation will be further enhanced by allowing for the eventuality when a certain treatment is discontinued, so that the virtual patient will revert to pre-treatment behavior in a realistic way.
We have added this functionality for all implemented diseases.

Objective 3 .The quality and coverage of the natural language processing substrate for the user-MVP communication will be strongly enhanced by accounting for "unexpected" inputs from the user; clarification sub-dialogs will be introduced into the system.
The natural language processing (NLP) aspect of the project has included the following: (1) developing new algorithms for the semantic processing of syntactically ill-formed or "unexpected" input; 2) increasing the lexicon by about 300 words/phrases that were drawn from on-line texts about esophageal diseases.

C. Smart Image

C. 1. Smart Image: CT guided imaging

Objective 1. Dose reduction strategy: High dose-Low dose. (Algorithmic developments; phantom and animal model study; validation)
Extending the work earlier reported in the publication (Dandekar O, Walimbe V, Siddiqui K, Shekhar R: Image registration accuracy with low-dose CT: How low can we go? In Proceedings of 2006 IEEE International Symposium on Biomedical Imaging, p 502-505.) to a 30-subject study including a rigorous validation of the performance of image registration. This preliminary study showed acceptable registration between standard-dose CT (simulating a pre-operative image) and very low-dose CT acquired at 11 mAs (simulating intra-operative image). This study is being extended. It is expected that two more quarters will be needed to complete this goal that includes preparation of a manuscript for publication in the journal "Radiology".

Objective 2. Dose reduction strategy: Iterative reconstruction (algorithmic developments; testing with scanner; software parallelization; hardware parallelization)
No significant progress has been made on dose reduction strategies using iterative reconstruction.

Objective 3. High-speed non-rigid (hardware implementation; software parallelization; hardware parallelization)
We presented a design of accelerated nonrigid image registration (Dandekar O, Walimbe V, Shekhar R: Hardware implementation of hierarchical volume subdivision-based elastic registration. In Proceedings of the 28th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (IEEE EMBC 2006), p 1425-1428). Hardware implementation of this design has become and will continue in the upcoming quarters. Preprocessing of CT images, especially low-dose CT images, is a prerequisite for reliable image registration. Moreover, image preprocessing needs to be no slower than the nonrigid image registration for the overall process to operate at high speed. The image preprocessing study is nearing completion. We recently submitted a manuscript on this topic to the Journal of Real-Time Image Processing. The details of the manuscript are: Dandekar O, Castro-Pareja CR, Shekhar R,"FPGA-based Reconfigurable 3D Image Pre-processing for Image Guided Interventions."

Objective 4. Tracking and visualization (Integration of tracking system; tool tracking for metal artifacts; automatic segregation; visualization; integration of visualization with tracking)
We have developed the mechanism to track the laparoscope and other tools optically. Using this setup we have also collected data to develop metal artifact correction algorithms, which are underway. In our preliminary work, using manual segmentation we have shown visualization of the operative field from the CT data. These techniques continue to be further developed and refined. Two conference presentations, one at MMVR and another at SPIE Medical Imaging (International Society for Optical Engineering), are planned in the first quarter of 2007 on the development of continuous CT-guided minimally invasive surgery. The paper is authored by Shekhar R, Dandekar O, Kavic S, George I, Mezrich R, and Park A.

C. 2. Smart Image: Image Pipeline

Computing architecture.
In order to support on-demand construction and display of complex imagery, an efficient software/hardware architecture is the prerequisite. We acquired a number of PCs for research assistants and software development. For data acquisition and experimentation, we have evaluated several tracking devices and non-contact scanning devices. A large display has been set up for visualization. It uses special purpose rendering hardware to obtain the maximum performance. On the software front, we have purchased several commercial software packages. In addition, custom software has been developed to process the data for the smart-imaging pipeline. In essence, we now have the necessary software/hardware components to acquire, process, and visualize the images, probably up to the extent any commercial system can offer. The remaining software/hardware components are in development.

Mental workload and other cognitive measures.
Mental workload and other cognitive measures are crucial indicators of how “smart” a smart image really is. Misplaced information focus can be a distraction. Misleading data can create inaccurate mental models. We have designed an experiment to quantitatively measure the effect of user interface design on mental workload. A user study has been conducted. The results from the user study are yet to be analyzed.

Multi-modal registration.
Multi-modal registration remains an open problem requiring serious research and engineering. The problem in smart imagery is particularly hard due to the following: (a) global coordinate frame for multiple sensors over time is difficult to establish, (b) pre-operative data do not easily align with intra-operative information, and (c) Non-rigid deformation creates challenges. We have evaluated a number of approaches, including external tracking (using the data captured at UMM), vision based 3D tracking, and 2D tracking. Each one of them takes a great deal of time and effort to implement and evaluate. Our investigations indicate that either external tracking or vision-based 3D tracking is able to produce stable results, in particular over long sequences. 2D tracking is

more stable, but requires an additional 2D to 3D registration step. This is one of the problems we will focus on this quarter.

The visualization and control of imagery in time and space.
The function is crucial to the usefulness of smart imagery since this is the part the surgeons/operators are going to use. We have developed software to show enhanced scope view, either in 2D or 3D, provided that the scope video data are already registered with 3D volumetric scans.

D. TraumaPod

This report covers activities initiated by the University of Maryland School of Medicine in support of the DARPA Trauma Pod project. UM is a supporting member of the Trauma Pod team and is providing software, clinical advice and design expertise for the automated supply dispensing system and developing the resource management system sub-components of Trauma Pod including the creation of an automated Medical Encounter Record and supply inventory monitoring.

Our primary goal during the reporting period was to develop a *Context-Aware Perioperative Information System* that would have the capability to automate most of the perioperative support functions such as patient tracking, inventory management, clinical documentation during the current and future phases of Trauma Pod with the help of pervasive computing and semantic web technologies. During Phase I the system would drive the Medical Encounter Record component. The task included developing a high level reasoning information system that utilized low level data streaming from a number of sources including physiological monitoring, inventory usage, tool usage and machine messaging within the Trauma Pod robotic cell. It is intended that the system design also anticipate the inclusion of RFID as a data source and we have designed the system to be RDID ready in future phases of Trauma Pod as well as traditional ORs. The following modules were designed and/or developed: (1) An RFID Server, (2) Database Interface, (3) Basic RMS Emulator, and (4) Design for the Medical Encounter Record.

System Architecture: The context-aware system is designed as a 3-tier event detection system. Events at the lower levels are processed to infer high-level events. Data is collected from various sensors in the operating room and this data is processed to reconstruct the surgical context and infer the medically significant events.

Data Sources:

Patient Monitors: TraumaPod Phase I utilizes the LSTAT critical care gurney as the basic platform for patient transport and all operative events. We use data streams from these patient monitors to determine the state of the patient during the surgery.

Network Messaging within the Robotic Cell: The MER monitors all network messaging within and among the robotic cell components.

Voice Commands: The MER also monitors voices commands as a form of system messaging that may have a causal or proximal relationships to medically significant events.

Data Stream Management System: TelegraphCQ:
The patient monitoring systems and the RFID reader produces continuous streams of data that need to be processed and analyzed in real-time to detect events. Traditional database systems have been designed to manage finite data sets where client queries are processed immediately against data stored in tables. In applications that process continuous data streams, clients require long running continuous queries that are evaluated as data streams through the application.

To manage the continual query function and process the resulting data, we used a data stream management system, TelegraphCQ, developed at University of California, Berkeley to process the physiological and RFID data streams. Data from patient monitors and the RFID reader is pushed to the stream engine continuously. Queries over these data streams are specified over a time window

Analyzing Physiological Data: Physiological parameters reflect a patient's health status. Interpretation of physiological data to infer the patient's condition is a challenging problem.

We used fuzzy set theory to capture this uncertainty in medical data. Given a data value, the memberships functions determine the degree to which the value belongs to a particular set. The value of the membership varies between 0 and 1 where 1 implies absolute membership. The set points used to define the range of values varies with each patient. For example, the range of normal blood pressure for a hypotensive patient will be different than the range for a patient with normal blood pressure. In the current version of the system, the set point for each parameter is preset for a patient.

During the reporting period we implemented a design change requested by SRI for the inventory interface that refines the use of EPC codes for identifying supplies and tools and provides a common ontology for the supply chain locations and supply status.

E. Video-assisted coordination tools

In the past reporting period we focused on the overall goal of supporting surgical workflow through advanced, real-time status display systems. We have continued analyzing data from our two existing sites, and published results as indicated. We plan to continue data collection in a third location—the post-anesthesia care unit. The data collection and analysis is contingent upon installation of the coordination display in the location in question.

Key Research Accomplishments

A. Informatics

Informatics subgroup 2. Operating Room Glitch Analysis (OGA)

- Developed several clinical scenarios with key stakeholders and subject matter experts.

- Set up database system and server to be in compliance with all HIPAA and security practices associated with PHI.

- Conducted literature review of peer reviewed publications on operating room activity analysis

Informatics subgroup 3. Context Aware Surgical Training (CAST)

- The CAST project is recently underway. No key research accomplishments to reported in this reporting period.

Informatics subgroup 4. Operating Room Clutter (ORC)

- Developed and piloted a set of preliminary measures

- Established and incorporated into the coding scheme a set of ergonomic problems associated with cords and cables used during minimal invasive procedures

- Piloted case analysis techniques to determine the effectiveness of coding schemes.

- Collected and cataloged literature on surgical ergonomics, operating room layout, and analysis of operations-related issues in OR management.

Informatics subgroup 5. Improving Perioperative Communications (IPC)

The progress in the IPC section has been impeded by the departure of a key personnel resource from the University of Maryland Medical Center, who provided both access to and information regarding operational data. While the IPC activities have continued and will continue with alternate personnel, the personnel change has slowed progress.

- Established high-level process mapping of communication, and preliminary needs assessment to determine key players and stakeholders.

- Analyzed information requirements for information system. (The sufficiency of this analysis is being examined.)

- Developed initial revision of an improved scripting for the scheduling process phone line.

- Analyzed available operational data to reveal deficiencies in the specificity of the currently-available data. Specifically, delay codes available in the current information system do not provide sufficient detail in their current format. Delay annotations, written in free text, are sometimes available, but these data cannot be analyzed using automated reporting systems standard data analysis techniques. Work being carried out in the

B. Simulation

- Added three new diseases to the simulation capabilities of the Maryland Virtual Patient (MVP): Zenker's diverticulum, scleroderma-esophagus, and LERD (the latter is partially implemented).

- Added this functionality for all implemented diseases.

The natural language processing (NLP) aspect of the project has included the following:

- Developed new algorithms for the semantic processing of syntactically ill-formed or "unexpected" input;

- Increased the lexicon by about 300 words/phrases that were drawn from on-line texts about esophageal diseases.

The simulation group has additionally accomplished the following:

- Created several new versions of the MVP system, in preparation for various demos. Each of the demonstration systems included enhanced functionality, improved look-and-feel, or furthered debugging of knowledge.

- The MVP system was made into a readily downloadable application that the physicians now have easy access to, for testing and knowledge debugging purposes.

- We began a concerted effort to shift the system from its current rapid prototyping level to a more stable, well-engineered environment.

C. Smart Image

C. 1. Smart Image: CT guided imaging

- Extended the work earlier reported to a 30-subject study including a rigorous validation of the performance of image registration.

- Presented a design of accelerated nonrigid image registration (Dandekar O, Walimbe V, Shekhar R: Hardware implementation of hierarchical volume subdivision-based elastic registration. In Proceedings of the 28th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (IEEE EMBC 2006), p 1425-1428).

- Hardware implementation of this design has been realized and will continue in the upcoming quarters. The image preprocessing study is nearing completion.

- Submitted a manuscript on this topic to the Journal of Real-Time Image Processing. The details of the manuscript are: Dandekar O, Castro-Pareja CR, Shekhar R,"FPGA-based Reconfigurable 3D Image Pre-processing for Image Guided Interventions."

- We have developed the mechanism to track the laparoscope and other tools optically.

- Collected data to develop metal artifact correction algorithms, which are underway.

- Shown visualization of the operative field from the CT data.

- Presented two conference presentations, one at MMVR and another at SPIE Medical Imaging (International Society for Optical Engineering), are planned in the first quarter of 2007 on the development of continuous CT-guided minimally invasive surgery. The paper is authored by Shekhar R, Dandekar O, Kavic S, George I, Mezrich R, and Park A.

C. 2. Smart Image: Image Pipeline

Computing architecture.

- Acquired a number of PCs for research assistants and software development. For data acquisition and experimentation.

- Evaluated several tracking devices and non-contact scanning devices.

- Set up a large display for visualization using special purpose rendering hardware to obtain the maximum performance

- Purchased several commercial software packages, and developed custom software to process the data for the smart-imaging pipeline. We now have the necessary software/hardware components to acquire, process, and visualize the images, probably up to the extent any commercial system can offer.

Mental workload and other cognitive measures.

- Designed an experiment to quantitatively measure the effect of user interface design on mental workload. Conducted a user study of mental workload measures.

Multi-modal registration.

- Evaluated a number of approaches, including external tracking (using the data captured at UMM), vision based 3D tracking, and 2D tracking.

The visualization and control of imagery in time and space.

- Developed software to show enhanced scope view, either in 2D or 3D, provided that the scope video data are already registered with 3D volumetric scans.

D. Trauma Pod

- Developed a high level reasoning information system that utilized low level data streaming from a number of sources including physiological monitoring, inventory usage, tool usage and machine messaging within the Trauma Pod robotic cell.

- Designed a 3-tier System Architecture for event detection in the context-aware to reconstruct the surgical context and infer the medically significant events.

Data Sources:

- Adapted the Medical encounter record to collect information from the embedded patient monitoring system on the LSTAT to the patients critical physiological parameters.

- Used data streams from these patient monitors to determine the state of the patient during the surgery.

Data Stream Management System: TelegraphCQ:

- Processed the physiological and RFID data streams using data stream management system TelegraphCQ, developed at University of California, Berkeley to manage the continual query function.

- Developed a means to simultaneously read several data sources that can connect to the stream-processing engine.

Analyzing Physiological Data:

- Used fuzzy set theory to capture "uncertainty" in medical data.

- Implemented a design change requested by SRI for the inventory interface that refines the use of EPC codes for identifying supplies and tools and provides a common ontology for the supply chain locations and supply status.

E. Video-assisted coordination tools

- Published four articles (see Reportable Outcomes).
- Prepared one article that is submitted for publication and under review.

Reportable Outcomes

A. Informatics

Agarwal, S., Joshi, A., Finin, T. and Yesha, Y. A Pervasive Computing System for the Operating Room of the Future, technical report TR-CS-06-xx, Computer Science and Electrical Engineering, University of Maryland, Baltimore County, September 2006. http://ebiquity.umbc.edu/paper/html/id/339/

Agarwal, S., Context-Aware System to Create Electronic Medical Encounter Records, M.S. thesis, Computer Science and Electrical Engineering, University of Maryland, Baltimore County, May 2006.
http://ebiquity.umbc.edu/paper/html/id/338/

Agarwal, S., Joshi, A., Finin, T., Ganous, T. and Yesha, Y. Context-Aware System to Create Electronic Medical Records, technical report TR-CS-06-05, Computer Science and Electrical Engineering, University of Maryland, Baltimore County, July 2006. http://ebiquity.umbc.edu/paper/html/id/312/

Vartak, N., Protecting the privacy of RFID tags. (M.S. thesis) Computer Science and Electrical Engineering, University of Maryland, Baltimore County, May 2006.

Publications under review
Vartak, N. Protecting the privacy of RFID tags, September 2006.

Agarwal, S., Joshi, A., Finin, T. and Yesha, Y. A Pervasive Computing System for the Operating Room of the Future, September 2006.

Publications in preparation:
Agarwal, S., Joshi, A., Finin, T. and Yesha, Y. Pervasive Computing in the Operating Room of the Future, to be submitted to Surgical Innovations, January 2006.

Grants pending

Context Awareness and Semantics in Pervasive Computing, National Science Foundation, 2007-2010, $449K, PI: A. Joshi, CO-PIs: T. Finin and Y. Yesha. Submitted December 6, 2006.
This proposal seeks funds to build a scalable software architecture to support intelligent pervasive information systems for an ORF like environment.

A Semantic Framework for Policy Specification and Enforcement in a Need to Share Environment, National Science Foundation, National Science Foundation, 2007-2010, $900K, PI: T. Finin CO-PIs: A. Joshi and Y. Yesha. Submitted January 8, 2007.
This will be done as a collaborative proposal with the University of Texas, Dallas and GMU. The total amount will be about $2M. A major use case and application testbed will be enforcing security and privacy policies in pervasive healthcare information systems like the ones we will be developing in ORF.

B. Simulation

Jarrell B, Nirenburg S, McShane M, Fantry G, Beale S, Mallott, D, Raczek J: “Simulation for Teaching Decision Making in Medicine: The Next Step”, In Proceedings of Medicine Meets Virtual Reality : IGT-Registration & Navigation, 2007.

C. Smart Image

Smart Image: CT guided imaging

Dandekar O, Walimbe V, Siddiqui K, Shekhar R: Image registration accuracy with low-dose CT: How low can we go? In Proceedings of 2006 IEEE International Symposium on Biomedical Imaging, p 502-505.

Dandekar O, Walimbe V, Shekhar R: Hardware implementation of hierarchical volume subdivision-based elastic registration. In Proceedings of the 28th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (IEEE EMBC 2006), p 1425-1428.

Dandekar O, Castro-Pareja CR, Shekhar R,"FPGA-based Reconfigurable 3D Image Pre-processing for Image Guided Interventions. Journal of Real-Time Image Processing. (Submitted)

Shekhar R, Dandekar O, Kavic S, George I, Mezrich R, Park A, "Development of continuous CT-guided minimally invasive surgery," In Proceedings of SPIE Medical Imaging 2007: Visualization and Image-Guided Procedures, vol. 6509, pp. 65090D, 2007

Shekhar R, Dandekar O, Kavic S, George I, Mezrich R, Park A, "Development of continuous CT-guided minimally invasive surgery," In Proceedings of Medicine Meets Virtual Reality: IGT-Registration & Navigation, 2007

D. Trauma Pod

- No publications at this time.

E. Video-assisted coordination tools

Publications:

Hu P, Xiao Y, Ho D, Mackenzie CF, Hu H, Voigt R, Martz D. Advanced Visualization Platform for Surgical Operating Room Coordination: Distributed Video Board System. Surgical Innovation. 13(2):129-135. 2006

Wasei M, Xiao Y, Wieringa P, Strader M , Hu P, Mackenzie C. Visualization of Uncertainty to Support Collaborative Trajectory Management in Hospital Care. 2006 Conference of International Ergonomics Association.3721-3726. 2006

Xiao Y, Strader M, Hu P, Wasei M, Wieringa P. Visualization Techniques for Collaborative Trajectory Management . ACM Conference on Human Factors in Computing Systems, pp.1547 - 1552. 2006

Xiao Y, Wasei M, Hu P, Wieringa P, Dexter F. Dynamic Management in Perioperative Processes: A Modeling and Visualization Paradigm. 12th IFAC Symposium on Information Control Problems in Manufacturing. (3)647-52. 2006

Paper under review:
Xiao Y, Dexter F, Hu P, Dutton RP. Usage of Distributed Displays of Operating Room Video when Real-Time Occupancy Status was Available. Anesthesia and Analgesia.

Conclusion

There are 5 major portions of this study: (a) OR informatics (b) simulation research, and (c) smart image, (d) TraumaPod, and (e) Video-assisted coordination tools.

The purpose of the OR informatics program is to develop, test, and deploy technologies to collect real-time data about key tasks and process elements in clinical operating rooms. We have established testbeds of activities in both simulated and operational environments. We are currently performing tests of the hardware, refining software, and applying lessons learned to hospital operational functions.

The objective of Simulation research is to create a system where a user can interact with a virtual human model in cognitive simulation and have the virtual human respond appropriately to user queries and interventions in clinical situations, with a focus on cognitive decision making and judgment. We have made significant strides toward realizing these goals. The MVP simulation functions well for esophageal disorders, and is continuing to expand the repertoire of diseases that are in the simulation model.

The objective of smart image is use real-time 3D ultrasonography and 40-slice high-frame-rate computed tomography (CT) for intraoperative imaging to volume rendered anatomy from the perspective of the endoscope. We are combining CT and Ultrasound to overlay image and data to enhance the performance of surgeons-in-training. We have carried out animate model testing of the image registration with great success. We continue to refine and expand our capability through hardware and software refinement.

The TraumaPod project. The University of Maryland is a supporting member of the DARPA-sponsored TraumaPod project, and is providing software, clinical advice, and design expertise for the automated supply dispensing system, and developing the resource-management system subcomponents of the TraumaPod Robotic Surgery System.

The Video-assisted coordination tools has developed and deployed an operational testbed to assist in coordination of two operating room suites with two prototype systems. Feedback from the operations of these prototypes is being applied to the design of a third prototype, to be deployed in the future.

References

1. Agarwal, S., Context-Aware System to Create Electronic Medical Encounter Records, M.S. thesis, Computer Science and Electrical Engineering, University of Maryland, Baltimore County, May 2006. http://ebiquity.umbc.edu/paper/html/id/338/

2. Agarwal, S., Joshi, A., Finin, T. and Yesha, Y. A Pervasive Computing System for the Operating Room of the Future, technical report TR-CS-06-xx, Computer Science and Electrical Engineering, University of Maryland, Baltimore County, September 2006. http://ebiquity.umbc.edu/paper/html/id/339/

3. Agarwal, S., Joshi, A., Finin, T. and Yesha, Y. Pervasive Computing in the Operating Room of the Future, to be submitted to Surgical Innovations, January 2006.

4. Agarwal, S., Joshi, A., Finin, T., Ganous, T. and Yesha, Y. Context-Aware System to Create Electronic Medical Records, technical report TR-CS-06-05, Computer Science and Electrical Engineering, University of Maryland, Baltimore County, July 2006. http://ebiquity.umbc.edu/paper/html/id/312/

5. Dandekar O, Walimbe V, Shekhar R: Hardware implementation of hierarchical volume subdivision-based elastic registration. In Proceedings of the 28th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (IEEE EMBC 2006), p 1425-1428.

6. Dandekar O, Walimbe V, Siddiqui K, Shekhar R: Image registration accuracy with low-dose CT: How low can we go? In Proceedings of 2006 IEEE International Symposium on Biomedical Imaging, p 502-505.

7. Hu P, Xiao Y, Ho D, Mackenzie CF, Hu H, Voigt R, Martz D. Advanced Visualization Platform for Surgical Operating Room Coordination: Distributed Video Board System. Surgical Innovation. 13(2):129-135. 2006

8. Jarrell B, Nirenburg S, McShane M, Fantry G, Beale S, Mallott, D, Raczek J: “Simulation for Teaching Decision Making in Medicine: The Next Step”, In Proceedings of Medicine Meets Virtual Reality : IGT-Registration & Navigation, 2007.

9. Nagy P, George I, Bernstein W, Caban J, Klein R, Mezrich R, Park A. Radio frequency identification systems technology in the surgical setting. *Surg Innov* 2006; 13(1):61-7.

10. Shekhar R, Dandekar O, Kavic S, George I, Mezrich R, Park A, "Development of continuous CT-guided minimally invasive surgery," In Proceedings of SPIE Medical Imaging 2007: Visualization and Image-Guided Procedures, vol. 6509, pp. 65090D, 2007.

11. Shekhar R, Dandekar O, Kavic S, George I, Mezrich R, Park A, "Development of continuous CT-guided minimally invasive surgery," In Proceedings of Medicine Meets Virtual Reality : IGT-Registration & Navigation, 2007.

12. Vartak, N. Protecting the privacy of RFID tags, M.S. thesis, Computer Science and Electrical Engineering, University of Maryland, Baltimore County, May 2006.

13. Wasei M, Xiao Y, Wieringa P, Strader M , Hu P, Mackenzie C. Visualization of Uncertainty to Support Collaborative Trajectory Management in Hospital Care. 2006 Conference of International Ergonomics Association.3721-3726. 2006.

14. Xiao Y, Strader M, Hu P, Wasei M, Wieringa P. Visualization Techniques for Collaborative Trajectory Management . ACM Conference on Human Factors in Computing Systems, pp.1547 - 1552. 2006

15. Xiao Y, Wasei M, Hu P, Wieringa P, Dexter F. Dynamic Management in Perioperative Processes: A Modeling and Visualization Paradigm. 12th IFAC Symposium on Information Control Problems in Manufacturing. (3)647-52. 2006

Paper under review:

1. Dandekar O, Castro-Pareja CR, Shekhar R,"FPGA-based Reconfigurable 3D Image Pre-processing for Image Guided Interventions. Journal of Real-Time Image Processing. (Submitted)

2. Xiao Y, Dexter F, Hu P, Dutton RP. Usage of Distributed Displays of Operating Room Video when Real-Time Occupancy Status was Available. Anesthesia and Analgesia.

Appendices

1. Jarrell B, Nirenburg S, McShane M, Fantry G, Beale S, Mallott, D, Raczek J: “Simulation for Teaching Decision Making in Medicine: The Next Step”, In Proceedings of Medicine Meets Virtual Reality : IGT-Registration & Navigation, 2007.

2. Shekhar R, Dandekar O, Kavic S, George I, Mezrich R, Park A, "Development of continuous CT-guided minimally invasive surgery," In Proceedings of Medicine Meets Virtual Reality : IGT-Registration & Navigation, 2007.

3. Shekhar R, Dandekar O, Kavic S, George I, Mezrich R, Park A, "Development of continuous CT-guided minimally invasive surgery," In Proceedings of SPIE Medical Imaging 2007: Visualization and Image-Guided Procedures, vol. 6509, pp. 65090D, 2007.

4. Agarwal, S., Joshi, A., Finin, T. and Yesha, Y. A Pervasive Computing System for the Operating Room of the Future, technical report TR-CS-06-xx, Computer Science and Electrical Engineering, University of Maryland, Baltimore County, September 2006. http://ebiquity.umbc.edu/paper/html/id/339/

Simulation for Teaching Decision Making In Medicine: The Next Step

Bruce JARRELL[a], MD, Sergei NIRENBURG[b], PhD, Marjorie MCSHANE[b], PhD, George FANTRY[a], MD, Stephen BEALE[b], PhD, David MALLOTT[a], MD, John RACZEK[a]
[a] *University of Maryland School of Medicine*
[b] *University of Maryland Baltimore County*

Abstract. This article describes the Maryland Virtual Patient project, which is developing realistic cognitive models of virtual patients for the training of medical students.

Keywords. cognitive simulation, virtual patient, medical education

Introduction

Clinical decision making skills are developed through practice on live patients. We train our physicians using the mix of live patients available at the time of training and trust that the knowledge acquired by managing this cohort of patients will be sufficient to generalize to all patients. Even though this is a tried and true method, it has many drawbacks related to patient safety, lack of objective measures for competence, and the inconsistency of the learning experience with respect to types of diseases, variations in the presentation or course of the disease, and individual patient differences. We need more opportunities to expose our trainees to sufficient patient scenarios in order to foster mastery of the complex knowledge needed daily by practicing physicians.

Computer-based simulation is one way to address the shortcomings of current clinical training practices. For simulation to be effective, it must expose the student to virtual patients that demonstrate sophisticated, realistic behaviors; it must allow open-ended patient investigation by the student (learning through self-discovery); and it must provide each student with a population of patients suffering from a given disease, with each patient displaying clinically relevant variations on the disease theme. Such variations might involve the path or speed of disease progression, the profile and severity of symptoms, responses to treatments, and secondary diseases or disorders that affect treatment choices. If each student could independently manage the care of many such patients – especially in a context in which trial and error learning carried no risk – we hypothesize that the decision making skills of each student would develop faster than with traditional training methods alone. In the Maryland Virtual Patient (MVP)[1,2] project we are developing a simulation and tutoring environment to test this hypothesis.

[1] Patent pending.
[2] This research was supported by Department of Defense grant #17-03-2-001.

Before launching full-scale work on this project, exploratory observational exercises were conducted with medical students at the University of Maryland School of Medicine to understand the specifications for effective interaction with a simulated patient [1].[3] In the exercises, the students managed several structured patients in electronic and manual simulations. All the exercises employed patient management problems used routinely in teaching and focused on high-level decision-making, such as the proposal and proof of an inference or the substantiation of an intervention. The most notable observations from this and a follow-up study of simulation for medical training were [1]:

- The simulation must accommodate trial and error patient management with multiple clinically plausible pathways to a solution.
- Changes in patient anatomy and physiology resulting from user action or disease processes over time must result in a consistent appropriate alteration of the state of the patient.

  The representation of time-related patient activities is critical for successful simulation, including allowing the user to "advance the clock" to the next phase of patient management.

In addition to these capabilities, the following are being incorporated into the MVP environment: chronic and acute disorders; simple and complex diseases; knowledge about well-understood and poorly-understood disease processes; knowledge spanning all levels, from gene to organism to population; complications of diseases and treatment modalities; and automatic tutoring. Among the most important conceptual aspects of MVP simulation is automaticity, which refers to the fact that the state of an MVP changes in a realistic way over time and in response to internal (physiological and pathological) and external (clinical and behavioral) stimuli.

Elicitation and Encoding of Knowledge

The MVP project centers on an ontology-based model of the physiological and cognitive processes affecting the virtual patient. We encode knowledge about biophysical functions that have clinical relevance in the maintenance of health, the production of disease, and the bidirectional transitions between these two states. When biomechanisms are known, they are modeled using causal chains. Where gaps exist in our knowledge of explicit biomechanisms, they are bridged in various ways – with non-biomechanistic knowledge from the literature, practical clinical knowledge, situational knowledge, observations, probabilistic methods, etc. This integration of implicit and explicit knowledge reflects precisely what a clinician employs when working with a patient. Further, the depth and granularity of this knowledge are determined by the demands of automatic function and realism. Thus, MVPs need not include every mechanism known to biology and clinical medicine.

Diseases are modeled as changes in key property values over time. For each disease, a set number of conceptual stages is established, and typical values (or ranges of values) for each property are associated with each stage. Values at the start or end of each stage are recorded explicitly, with values between stages being interpolated. The disease model includes a combination of fixed and variable features. For example,

[3] The exercises were conducted under IRB Exemption No. BJ-090103 for the project entitled "Computer Simulation as an Aid to Enhancing Medical Education."

although the number of stages for a given disease is fixed, the duration of each stage is variable; similarly, although the values for *some* physiological properties undergo fixed changes across patients, the values for other physiological properties are variable within a specified range. Therefore, on the one hand, each disease model is sufficiently constrained so that MVPs suffering from the disease must show appropriate physiological manifestations of it, while on the other hand, each disease model is sufficiently flexible to permit instances of MVPs to differ in clinically relevant ways, as selected by the author of each MVP instance.

Once an approach to modeling a given disease has been devised and all requisite details have been elicited, the disease-related events and their participants are encoded in ontologically-grounded scripts written in the metalanguage employed in the OntoSem environment.[4] Scripts represent typical sequences of events and their causal and temporal relationships. In other words, they encode how individual events hold well-defined places in routine, typical sequences of events that happen in the world, with a well-specified set of objects filling different roles throughout that sequence. For example, if the event is swallowing, there is only one animate participant (the swallower), but many other objects play necessary roles: various nerves and muscles act as instruments of peristalsis; the swallowed bolus is the theme of peristalsis-driven motion events; the stomach is the final destination of the bolus, and so on. Scripts normally contain subscripts and can be more or less fine-grained depending on the goals of the given simulation. Within the MVP project we have developed both domain scripts and workflow scripts. Domain scripts describe basic physiology, disease progression and responses to treatments, whereas workflow scripts model the way an expert physician would handle a case, thus forming the knowledge substrate for automatic tutoring.

The OntoSem ontology differs from others not only in its inclusion of scripts, but also in its rich inventory of properties, both attributes and relations (most other ontologies, e.g., UMLS [3], are actually hierarchical word nets rather than knowledge-rich ontologies). As we expand the OntoSem general-purpose ontology into the medical domain, we are incorporating, where possible, the terminology used by the Foundational Model of Anatomy [4].

Reasoning with Knowledge

MVPs are modeled as "double agents" with both physiological and cognitive functions. Physiologically, the state of an MVP changes in response to internal pathophysiological stimuli and external stimuli, the latter initiated either by the patient or the trainee. Cognitively, the MVP can communicate with trainees about current symptoms, lifestyle, history, adherence to prescribed treatments, etc. Structured knowledge in the disease model acts as input to the simulation engine. The simulation can run in clinical mode, where patient symptoms and physiology are known only through questions and diagnostic tests, or in omniscient mode, in which all patient properties can be monitored throughout the simulation. As an example of automaticity in response to external interventions, students are permitted to prescribe any treatment available in the system at any time, with the MVP responding accordingly. If, for

[4] OntoSem is the implementation of the theory of Ontological Semantics, a theory originally developed for knowledge-rich text processing (Nirenburg and Rakin 2004).

example, the student launches an inappropriate treatment, the MVP's state may or may not change, but certainly will not produce the intended result. Upon recognizing this undesirable result, the student can attempt to recover from the mistake, for example by withdrawing the treatment or introducing a different one. The effect of recovery attempts are interpreted relative to encoded knowledge and the current state of the MVP. The system does not exhaustively list all permutations of paths a trainee could take and all consequential responses of the MVP; instead, it relies on ontologically-grounded descriptions of basic physiology, disease processes, effects of treatments, and so on, so that the state of a given MVP at a given time will, quite literally, fall out of the underlying model.

Authoring Instances of MVPs

A cornerstone in creating a realistic MVP environment is providing for wide variation among instances of MVPs with a given disease. That is, the basic model of a disease includes all relevant tracks (i.e., paths of progression), and each track provides many choice points that differentiate cases. Among the many tasks carried out by the author of a disease model is the selection of properties to be tracked, their ranges of values, and the defaults for those values. Such a disease model is then concretized into a given patient instance by an instance author, who is typically a physician-teacher or disease specialist. The instance author determines the MVP's basic physiological properties, relevant lifestyle factors, the rate of progression of the disease, which path the disease takes at all possible furcations, the specific symptom profile at given times, and so on. This process has been reduced to an electronic multiple-choice questionnaire that takes little time to complete. The simplicity of authoring patient instances derives from the care taken to create the basic model of the disease, including delineating exactly which property values are available for individual parameterization and which ones are fixed for all patients experiencing the given stage of a disease. As soon as the values (or defaults) for all relevant properties are chosen, the patient instance is available for use.

Results

Knowledge Elicitation. The MVP project places significant demands on authors of disease models to render complex, multi-scale functions in a form that can be implemented computationally. The knowledge elicitation process is a collaboration between the model author and a knowledge engineer, who mediates between the physician and the programmer. Physicians must distill their extensive and tightly coupled physiological and clinical knowledge into the most relevant subset, and express it in the most concrete of terms. Not infrequently, they are also called upon to hypothesize about the unknowable, like the state of a patient experiencing a pre-clinical stage of disease, or the state of a patient after an effective treatment that is never, in real life, followed up by objective tests. Such hypotheses reflect the mental models of given experts, which might differ in subtle ways from those of other experts. However, such differences, we would suggest, have little bearing on the ultimate goal of this enterprise: to create MVPs whose behavior is sufficiently life-like to further specific teaching goals.

Disease Simulation. We chose to initially model esophageal disease because the esophagus is a relatively uncomplicated organ and because one of the symptoms of esophageal disease, chest pain, can cause significant diagnostic dilemmas with cardiac disease. Knowledge about the normal and abnormal anatomy and physiology of the esophagus was elicited from model authors and recorded. The two common mechanisms for gastroesophageal reflux disease (GERD) – a decreased Lower Esophageal Sphincter Pressure (LESP) and Transient LES Relaxation (TLESR) – and all relevant clinical forms of GERD were modeled. The latter included: non-erosive GERD; GERD with erosive esophagitis, stricture, Barrett's metaplasia and adenocarcinoma; and proximal GERD. In addition, diseases potentially associated with GERD, including scleroderma, Zenker's Diverticulum and achalasia, were modeled. For achalasia and scleroderma, which are diseases with a poorly understood esophageal pathophysiology, property values of the MVP change as a function of passing time, since the disease natural history can only be clinically observed rather than explained using causal chains. By contrast, for GERD, which has a well understood pathophysiology, the disease model is driven by causal chains that reflect current biomedical thinking.

Causal chain modeling is a particularly potent strategy that allows expanded opportunities for automatic function in virtual patients:

- A new disease can be generated as a side-effect of another disease: for example GERD is automatically initiated in any patient whose LESP drops below 10, which can occur due to scleroderma or after a successful surgical intervention for achalasia.
- The rate of progression of a disease can be automatically determined: for example, a patient with an LESP of 0 (after a successful Heller myotomy) will have a faster progression of GERD than a patient with an LESP of 9 (after a successful pneumatic dilation).
- The effects of interventions can be automatically determined: for example, whether GERD is progressing or healing is determined by the daily total time in acid reflux (TTAR). TTAR is determined by total time in reflux – which can be altered in some patients by changes in lifestyle, and by the acidity of the refluxed substance – which can be affected medication.

Upon testing, the MVP has functioned accurately, including reasoning effectively when responding to both expected and unexpected user actions.

Discussion

We have designed a simulation system that has demonstrated complex, automatic behavior. Thus far, it has a limited repertoire of esophageal physiology, pathophysiology, and clinical management for common esophageal diseases. In spite of the limited repertoire, we believe that MVPs represent a conceptual leap in the computer modeling of humans in the continuum of health and disease. MVPs display realistic function in simulations where they can be observed, interacted with, and treated by students. Variability of selected parameters permits a wide variety of instances of virtual patients to be created from the same ontologically-grounded disease model. We have honed our approaches to knowledge elicitation, script writing and incorporating scripts into the simulation engine, facilitating the expansion of the system's purview to other organs and systems. We are currently modeling ischemic cardiac disease, which will give us the opportunity to present a wide variety of MVPs with chest pain from either cardiac or esophageal origin.

When the model of heart disease is completed, we will be positioned to test our hypothesis that trial and error management of MVPs can teach students the basics of clinical medicine as well as or better than bedside teaching or small group teaching. In fact, there is evidence that learning by working through computer-based scenarios can be very effective: for example, in the evaluation of the SHERLOCK II system, which teaches electronics troubleshooting, it was reported that technicians learned more from using this system for 24 hours than from 4 years of work in the field [5].

Two components are necessary for a trainee to learn clinical medicine: an inventory of patients showing clinically relevant variations of disease, and a tutor to guide the student (as necessary) and to validate that his or her success derives from accurate and sufficient knowledge. We have recently implemented the first version of the tutor for esophageal diseases.

Our work on tutoring has been informed by results from the CIRCSIM group, which has been pursuing automatic tutoring strategies for the diagnostics and treatment of the baroreceptor reflex. The current CIRCSIM-Tutor evolved from a system that offered students a dynamic mathematical model with no tutoring support into a system that offers tutoring without the dynamic mathematical model (results of certain scenarios were stored and are deemed sufficient for the given educational goals). The contrast with MVP is clear: for us, the autonomous functioning of MVPs is, and will remain, central, with tutoring being interpreted as a useful option alongside trial-and-error learning.

We are also currently working on incorporating natural language interaction into the system, with the current mode of interaction being menu-driven (that is, the trainee is presented with inventories of questions, diagnostic tests, treatments, hypotheses and diagnoses to choose from). The desire to incorporate natural language interaction into tutoring systems has been expressed by developers of many tutoring systems. Unlike others, however, our group has been working on knowledge-based natural language processing (NLP) for some twenty years. In fact, the OntoSem ontology, knowledge representation language, and many of the processors that are serving as a substrate for the MVP system were all originally developed for NLP applications. Therefore, we have confidence in our ability to incorporate natural language support into the MVP environment in the near term.

References

[1] D. Mallott, J. Raczek, C. Skinner, K. Jarrell, M. Shimko and B. Jarrell. 2005. A basis for electronic cognitive simulation: the heuristic patient. *Surgical Innovations* March, 12(1):43-9.

[2] S. Nirenburg and V. Raskin. 2004. *Ontological Semantics*. The MIT Press.

[3] O. Bodenreider. 2004. The Unified Medical Language System (UMLS): Integrating biomedical terminology. *Nucleic Acids Research* 32: 267-270.

[4] C. Rosse and J.L.V. Mejino. 2004. A reference ontology for bioinformatics: The Foundational Model of Anatomy. *Journal of Biomedical Informatics*.

[5] M. Evens and J. Michael. 2006. *One-on-One Tutoring by Humans and Computers*. New Jersey and London: Lawrence Erlbaum and Associates, Publishers.

Development of Continuous CT-Guided Minimally Invasive Surgery

Raj Shekhar, Omkar Dandekar, Steven Kavic, Ivan George, Reuben Mezrich, Adrian Park
University of Maryland, Baltimore; rshekhar@umm.edu

Minimally invasive surgeries performed under laparoscopic guidance lead to improved patient outcomes, less scarring and significantly faster patient recovery as compared to conventional open surgeries. Rigid endoscopes (laparoscopes) are used to visualize internal anatomy and guide laparoscopic surgeries. Laparoscopes, however, are limited in their visualization capability by their flat representation of three-dimensional (3D) anatomy and their ability to display only the most superficial surfaces. A surgeon is unable to see beneath visible surfaces, decreasing the precision of current-generation laparoscopic surgeries. Awareness of the 3D operative field is a long-standing need of laparoscopic surgeons that laparoscopes are fundamentally limited in meeting.

Our solution to this problem is to use continuous computed tomography (CT) of the operative field as a supplementary imaging tool to guide laparoscopic surgeries. 3D visualization of anatomical structures from CT data is common in diagnostic radiology. Moreover, it is possible to expose hidden structures or to see inside organs by "peeling off" outer layers by making corresponding voxels transparent. The recent emergence of 64-slice CT as well as its continuing evolution in speed and volumetric coverage makes it an ideal candidate for four-dimensional (3D space + time) intraoperative imaging. Cost and availability considerations and the ability to image across pneumoperitoneum (caused by CO_2 insuffulation) also favor CT.

Our initial attempts have focused on dose reduction and a preliminary demonstration of 3D visualization of the operative field using continuous CT. To minimize net radiation dose to the patient and the surgeon with the use of continuous CT, we have proposed a novel dose reduction strategy, in which we acquire a standard CT image preoperatively (following pneumoperitoneum) and scan the dynamic operative field using very low-dose CT once surgery begins. Using high-speed nonrigid 3D image registration (warping) techniques we have developed [1-2], we rapidly register the preoperative CT image to low-dose intraoperative CT images. Registered preoperative CT images, which match the intraoperative anatomy, are then substituted for the low-dose images, 3D rendered and presented to the surgeon.

Our simulation experiments, designed to measure the accuracy of nonrigid registration at various simulated doses, show that intraoperative tissue shifts can be tracked with an accuracy of 2 mm even at an x-ray tube current of 10 mAs (typical diagnostic imaging uses 200 mAs) [3]. This is equivalent to a 94% and 95% reduction in the surface and the deep tissue dose, respectively, indicating that continuous CT can provide safe and accurate surgical guidance.

In an early demonstration of CT-based intraoperative visualization, a swine, lying supine on the CT couch, was prepared for laparoscopic cholecystectomy as per the standard procedure. The laparoscope was tracked in space using an optical tracker, which was calibrated with the reference frame of the CT acquisition. The animation in Fig. 1 displays side-by-side the laparoscopic and CT-generated views of the liver and the surrounding anatomy. Note that the hepatic vessels hidden beneath the liver surface are visible in the CT view. Visualization of critical structures, such as the vasculature, is important before making surgical dissections.

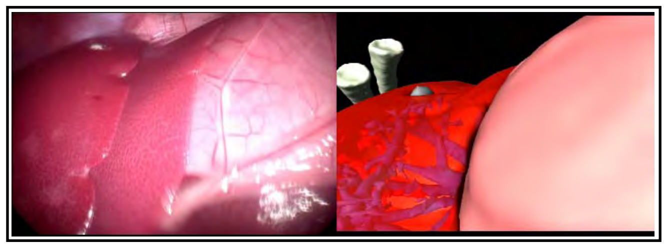

Figure 1 Laparoscopic display (left) versus the matched CT-generated view (right) [click for a movie]

Continuous low-dose CT scanning of the dynamic operative field without exceeding acceptable radiation exposure and using high-speed nonrigid registration to generate diagnostic quality CT images of the intraoperative anatomy will enable high-quality 3D visualization of operative field in a CT-equipped operating room (OR). We have successfully created 3D renderings from multi-slice CT corresponding to anatomy visible within the field of view of the laparoscope. We will extend this work to create augmented reality views as previously reported, albeit using instantaneous CT images. With additional developments, our research has the potential to help improve the precision of laparoscopic surgeries further, reduce complications, and expand the scope of minimally invasive surgeries to beyond its current 15% share of all surgeries.

References

1. Castro-Pareja CR, Jagadeesh JM, Shekhar R. FAIR: A hardware architecture for real-time 3-D image registration. *IEEE Trans Inf Technol Biomed* 2003;7(4):426-434.
2. Castro-Pareja CR, Shekhar R. Hardware acceleration of mutual information-based 3D image registration. *J Imaging Sci Technol* 2005; 49(2):105-113.
3. Dandekar O, Walimbe V, Siddiqui K, Shekhar R. Image registration accuracy with low-dose CT: How low can we go? In *Proceedings of 2006 IEEE ISBI*; p 502-505.

Development of continuous CT-guided minimally invasive surgery

Raj Shekhar[1,2,*], Omkar Dandekar[1,2], Steven Kavic[3], Ivan George[3], Reuben Mezrich[1], Adrian Park[3]

[1]Department of Diagnostic Radiology, University of Maryland, Baltimore, MD 21201, USA
[2]Department of Electrical & Computer Engineering, University of Maryland, College Park, MD 20742, USA
[3]Department of Surgery, University of Maryland, Baltimore, MD 21201, USA

ABSTRACT

Minimally invasive laparoscopic surgeries are known to lead to improved outcomes, less scarring, and significantly faster patient recovery as compared with conventional open invasive surgeries. Laparoscopes, used to visualize internal anatomy and guide laparoscopic surgeries, however, remain limited in visualization capability. Not only do they provide a relatively flat representation of the three-dimensional (3D) anatomy, they show only the exposed surfaces. A surgeon is thus unable to see inside a structure, which limits the precision of current-generation minimally invasive surgeries and is often a source of complications. To see inside a structure before dissecting it has been a long-standing need in minimally invasive laparoscopic surgeries, a need that laparoscopy is fundamentally limited in meeting. In this work we propose to use continuous computed tomography (CT) of the surgical field as a supplementary imaging tool to guide laparoscopic surgeries. The recent emergence of 64-slice CT and its continuing evolution make it an ideal candidate for four-dimensional (3D space + time) intraoperative imaging. We also propose a novel, elastic image registration-based technique to keep the net radiation dose within acceptable levels. We have successfully created 3D renderings from multislice CT corresponding to anatomy visible within the field of view of the laparoscope in a swine. These renderings show the underlying vasculature along with their latest intraoperative orientation. With additional developments, our research has the potential to help improve the precision of laparoscopic surgeries further, reduce complications, and expand the scope of minimally invasive surgeries.

Keywords: Image-guided surgery, laparoscopic surgery, augmented reality, elastic registration, 3D visualization

1. INTRODUCTION

Minimally invasive surgeries are a superior alternative to conventional open surgeries. In minimally invasive surgeries, the internal anatomy is accessed through a few small ports (holes) on the patient's skin rather than large incisions. The surgeon introduces the laparoscope (rigid endoscope) through one of the ports to illuminate the internal anatomy and uses the other ports to introduce surgical instruments. The region is often insufflated (filled) with CO_2 gas to make space for surgical manipulations and to provide access to the anatomy of interest. Minimally invasive surgeries performed under laparoscopic guidance have been shown to lead to improved outcomes, less scarring, and significantly faster patient recovery as compared to conventional open surgeries.[1] For certain surgical procedures, such as cholecystectomy (removal of gall bladder), minimally invasive surgery has become the standard of care.[2]

Despite the early success of minimally invasive surgeries, laparoscopes remain limited in visualization capability by their flat representation of three-dimensional (3D) anatomy and their ability to display only the most superficial surfaces. A surgeon is thus unable to see inside or around a structure, thus decreasing the precision of current-generation laparoscopic surgeries. Laparoscopic surgeons need awareness of the 3D operative field, especially visualization of the underlying blood vessels and other hidden structures.[3] Laparoscopes are fundamentally limited in providing this information and unable to meet this long-standing need.

---

[*] rshekhar@umm.edu; phone (410)706-8714; facsimile (410)706-8724

Augmented reality (AR) has been suggested as a solution to overcome this limitation of laparoscopy. One approach to AR has been the creation of 3D models of organs from preoperative magnetic resonance (MR) or computed tomography (CT) images acquired days to weeks before the surgery.[4, 5] Such 3D visualization of anatomical structures from CT or MR imaging data is common in diagnostic radiology. Moreover, it is possible to expose hidden structures or to see inside organs by "peeling off" outer layers by making corresponding voxels transparent. These models have been subsequently rendered with exquisite detail and superimposed on real-time laparoscopics display for AR. Although this approach has shown the strength of combining CT or MR image-based visualization with laparoscopy, its accuracy remains suspect. The 3D organ models derived from preoperative CT or MR imaging are not current and do not represent the intraoperative anatomy that almost invariably will deform between the time of preoperative imaging and the surgical procedure.

We propose improving intraoperative visualization during laparoscopic surgeries through AR that uses 3D renderings of the anatomy scanned with live, intraoperative CT. Superimposition of such 3D views based on instantaneously acquired CT on the laparoscopic view after accounting for proper alignment has the potential to reveal hidden structures accurately and thereby assist the laparoscopic surgeon. Although computationally and practically more challenging, this approach does not suffer from the limitations of previously reported AR efforts. With the advent of 64-slice CT scanners, continuous intraoperative volumetric CT at high frame rates is becoming possible. The continual trend toward more slices (hence, volumetric coverage per rotation) and higher frame rate will make CT even more suitable for this surgical imaging task.

Radiation exposure to the patient and the surgeon is a concern with the use of continuous CT as proposed here. In this article, we not only present our preliminary results showing AR visualization using intraoperative CT but also describe a strategy to reduce the radiation dose based on registration of pre- and intraoperative CT. We conclude with a discussion of our results, strengths of our proposed strategy, and future directions of our research.

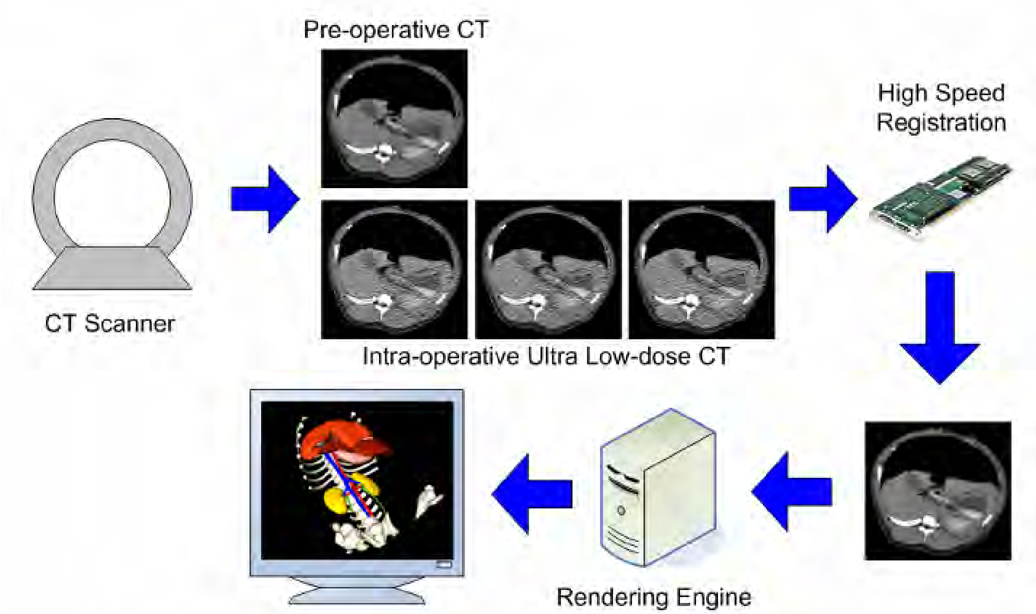

Figure 1 Flow diagram showing the CT protocol and various other steps involved in the proposed CT-guided laparoscopic surgery. Note that elastic image registration between standard-dose preoperative CT and ultra low-dose intraoperative CT helps reduce radiation exposure. A high-speed image registration engine in the final implementation will allow continuous intraoperative 3D visualization.

2. METHODS

1.1. Imaging protocol and dose reduction strategy

Figure 1 schematically describes our imaging protocol for the proposed CT-guided laparoscopic surgery and the concept behind our novel dose reduction strategy. After the surgery subject has been prepared (preparation includes insufflation) and immediately before surgery begins, we perform a contrast-enhanced CT scan at the standard diagnostic dose. The use of the contrast agent ensures that the desired blood vessels are highlighted in the CT. We call this initial 3D scan the preoperative CT scan. As soon as the surgery begins, we scan the dynamic operative field repeatedly with CT again, although this time the CT scanner is operated at a much lower dose. We refer to these subsequent nondiagnostic scans as intraoperative CT scans. Intraoperative CT scans are not contrast enhanced because of the short-acting nature of the CT contrast agents and the fact that these agents cannot be administered repeatedly without stress and potential harm to the kidneys and other critical organs.

Our next step is to elastically register pre- and intraoperative CT scans, which allows us to warp the diagnostic-quality preoperative scan in such a way that it matches the intraoperative anatomy. The warped preoperative CT scan, which has superior image quality, is then substituted for the intraoperative CT scan. This scan is subsequently rendered and superimposed on the corresponding laparoscopic view for AR. By repeating this process for each intraoperative CT scan several times per second, our approach can provide an accurate and up-to-date AR visualization throughout the surgery.

The dose reduction results from the use of ultra low-dose CT intraoperatively. Diagnostic CT scans are typically acquired at an x-ray tube voltage of 120 kVp and a tube current of 200 mAs. Working with a commercial CT dose simulator and archived patient images, we showed previously that accurate elastic image registration can be achieved even when the tube current is lowered to 10 mAs, thus resulting in an approximately 20-fold reduction in the radiation dose.[6] The dose reduction was slightly less when working with a commercial CT scanner because of preset limits on minimum tube current setting.

1.2. CT-laparoscopy spatial correlation

We performed the proposed continuous CT-guided laparoscopic surgery in a CT room working with a 64-slice scanner (Brilliance 64, Philips Medical System, Highland Heights, OH). Figure 2 shows our experimental setup. One of the first requirements for this novel surgical approach and AR visualization is to establish a spatial correlation between the CT coordinate system and the coordinate system of the laparoscope. Once initialized, the coordinate system of the CT scanner remains fixed. Because the laparoscope is manually operated, its coordinate system moves with it. We achieved the necessary spatial correlation through the use of an optical tracker (Polaris Spectra, Northern Digital, Waterloo, Canada). Infrared markers were attached to the external end of the laparoscope. A PC (termed the control PC) controlled the optical tracker, which, by tracking the infrared markers, provided the coordinates of the laparoscope in its (optical tracker's) coordinate space. A purpose-built calibration device with infrared markers visible in CT as well as visible to the optical tracker's cameras helped determine the transformation between the CT and the optical tracker coordinate systems.

The control PC was also fitted with a video frame grabber to capture and digitize the laparoscopic video. The control PC also synchronized the tracking and the video data. Temporal synchronization of laparoscopic data with the CT images is manual in our current system. Actions such as sudden movement of the laparoscope provided visual cues to synchronize CT and laparoscopy data.

1.3. Experimental protocol

In this preliminary study, our goal was to (1) prove the feasibility of accurate elastic registration between diagnostic preoperative CT and ultra low-dose intraoperative CT scans and the resulting radiation dose reduction, and (2) develop the necessary engineering tools to demonstrate AR visualization using warped preoperative CT. The focus was to achieve this goal for discrete time instants. By looping through these steps at a fast rate, the process can be made continuous in the future. Consequently, we followed the imaging protocol described in Figure 1, with the exception that continuous intraoperative CT was replaced with discrete single CT scanning. In addition, intraoperative CT and laparoscopy were performed sequentially while keeping the anatomic deformations to the minimum between these.

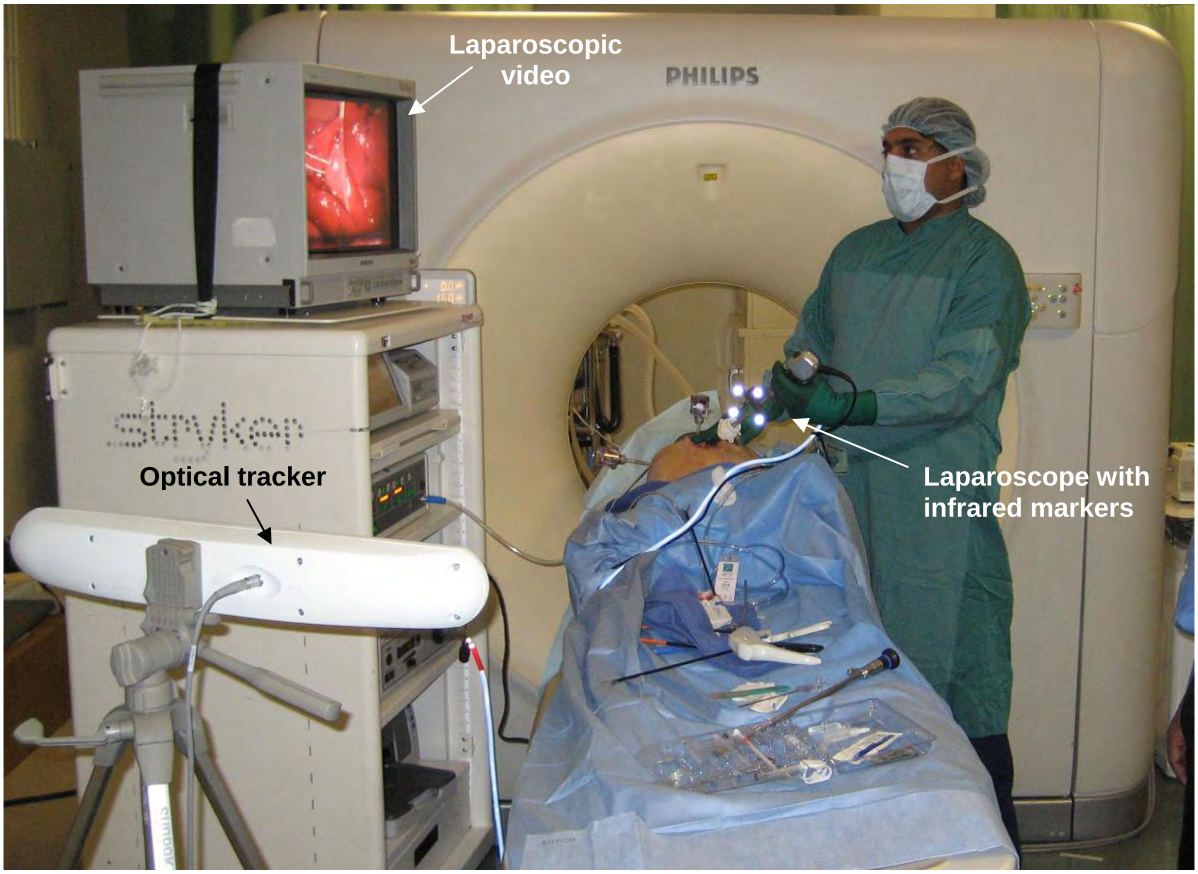

Figure 2 The setup for CT-guided minimally invasive laparoscopic surgery. The laparoscope is tracked using an optical tracker. The AR views, in this preliminary study, were created in an offline fashion.

We tested the proposed continuous CT-guided surgery approach on an anesthetized swine. Appropriate institutional approval was obtained for this animal protocol. Because we aim to demonstrate the feasibility of CT guidance for cholecystectomy (one of the most common minimally invasive surgeries) first, the particular organ of interest naturally was the gall bladder. The field of view for both CT and laparoscopy also included nearby organs such as the liver and the spleen. A preoperative CT scan with contrast-enhanced hepatic vasculature was acquired first at the standard dose per the protocol. A tube current setting of 100 mAs was used to account for preoperative CT to make adjustments for the smaller body weight of the animal. To simulate a surgical manipulation, the gall bladder was mobilized (i.e., separated from the liver and then elevated above the liver surface). Intraoperative CT scans were subsequently acquired at 4 different dose settings (100, 45, 30, and 15 mAs). Fifteen mAs was the lowest dose setting allowed by the scanner for the desired 1-mm CT. The x-ray tube voltage was kept fixed at 120 kVp.

1.4. Elastic image registration and AR visualization

The standard-dose preoperative CT scan and the low-dose intraoperative CT scans were registered as reported in our earlier dose simulator study using our own volume subdivision-based elastic image registration algorithm.[6,7] As before, the low-dose CT scans were preprocessed using an anisotropic diffusion filter. The registration was repeated for each of the four dose settings. The accuracy of registration was judged visually by evaluating the difference of warped preoperative CT and preprocessed intraoperative CT.

Independently, the organs and structures of interest (liver, gall bladder, and hepatic vasculature) in the preoperative CT scan were segmented using Amira 3D visualization software (Mercury Computers, Chelmsford, MA). The segmented structures were transformed according to the deformation field provided by elastic image registration with the intraoperative scan. The appropriately deformed segmented structures were surface rendered (again using Amira) from a series of viewing angles provided by the tracking data for the laparoscope. These rendered scenes were then superimposed with the corresponding laparoscopic video frames for AR visualization.

3. RESULTS

Our earlier simulation experiments, designed to measure the accuracy of nonrigid registration at various simulated doses, showed that intraoperative anatomical shifts can be tracked with an accuracy of 2 mm even at an x-ray tube current of 10 mAs (typical diagnostic imaging uses 200 mAs).[6] This is equivalent to a 94% and 95% reduction in the surface and deep tissue doses, respectively, indicating that continuous CT can provide safe and accurate surgical guidance. Figure 3 shows no visually noticeable difference in the performance of elastic image registration at high and low doses.

The same experiments were repeated for the swine data. In this case, the standard preoperative CT dose was lower (100 versus 200 mAs) to match the lower body weight of the animal compared with that of an adult human. In Figure 4, the pre- and intraoperative images are fused using a two-color scheme (red + blue channels for one image, green for another). The left panel shows the fused image before performing image registration. Bony structures such as the spine and ribs are clearly misaligned. The alignment improves after elastic image registration (middle and right panels). In this case, too, no visually noticeable difference in the registration quality is seen in images at the standard dose (100 mAs) or the lowest dose (15 mAs), indicating the potential of ultra low-dose intraoperative CT.

The procedures described above demonstrated the feasibility of substituting warped preoperative CT for intraoperative CT and reducing radiation exposure in the process. The warped preoperative CT is further used for high-quality 3D visualization and for creating a more accurate AR. Figure 5 displays side by side the laparoscopic and CT-generated views of the liver and the surrounding anatomy for a given time instant. Note that the hepatic vessels hidden beneath the liver surface are visible in the CT view. Laparoscopy cannot provide such information.

It is also important to note the benefit of registration for the visualization of the vasculature. Because the contrast agent cannot be used repeatedly, the vessels are not enhanced in intraoperative CT images. Even if the noise in the low-dose CT could be suppressed through a filtering operation, rendering intraoperative CT directly would not reveal the vasculature (see Figure 6, left panel). Only a few major vessels are visible here because of the residual CT contrast agent in them. These deep-seated vessels are not significant for a surgery like cholecystectomy. The contrast agent washes out fairly rapidly from the most peripheral vessels, which are of primary interest in most image-guided surgeries. Image registration allows use of the preoperative data for 3D visualization of the intraoperative anatomy while also retaining the vasculature information (see Figure 6, right panel).

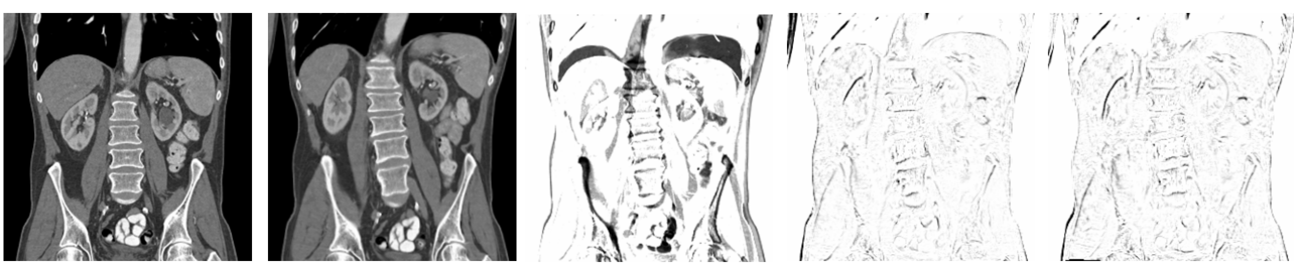

Standard CT (Pose 1) | Standard CT (Pose 2) | Starting Difference | 200 mAs (Standard dose) | 10 mAs (Ultra low dose)

Difference after elastic registration at standard dose and ultra low dose

Figure 3 Demonstration of elastic image registration accuracy for standard dose–low dose CT registration. The difference images suggest no visually noticeable difference in registration accuracy with dose. Quantitative data support this finding.

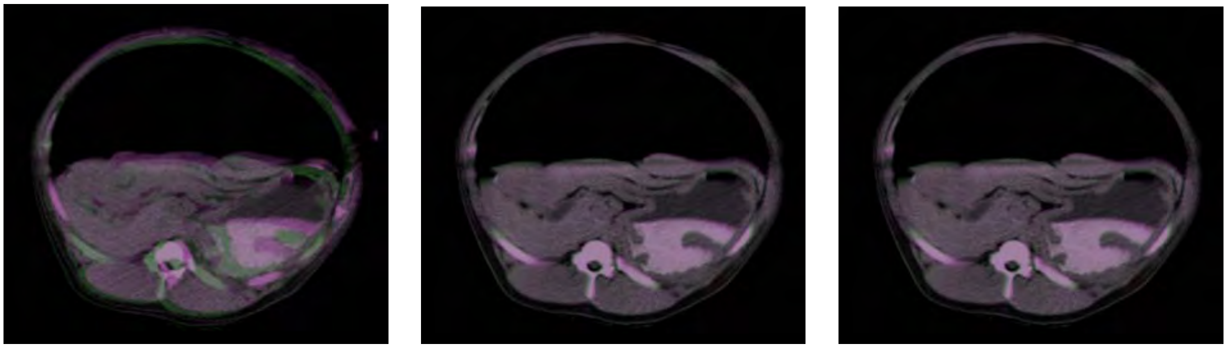

Before registration | After registration (100 mAs) | After registration (15 mAs)

Figure 4 Demonstration of elastic image registration accuracy for standard dose (100 mAs) – low dose (15 mAs) CT images of the swine. The difference images suggest no visually noticeable difference in registration accuracy even at a very low dose.

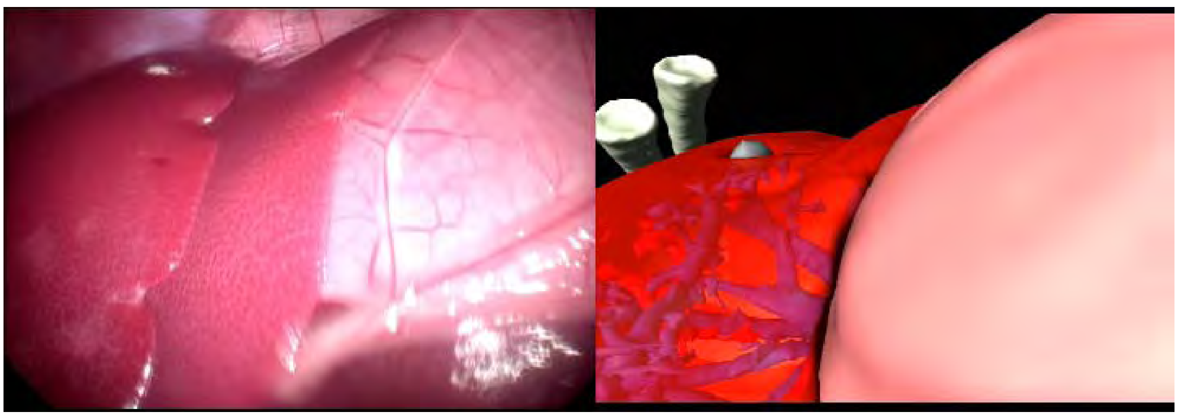

Figure 5 A comparison of spatially matching laparoscopic and CT-generated views. The CT-generated view (right panel) is capable of revealing the underlying vasculature, visualization of which is beneficial to laparoscopic surgeons.

Figure 6 Because of the inability to reuse contrast material, a direct rendering of the preoperative CT (left), cannot show the detailed vasculature. The rendered warped CT (right) shows the vasculature.

Finally, we created AR views by blending semitransparent laparoscopic and CT-generated views. An example of such AR visualization using the instantaneous volumetric CT scan of the intraoperative anatomy is shown in Figure 7. Such an AR view preserves the surface texture information and optical depth cues from the laparoscopy while also exposing the underlying vasculature accurately.

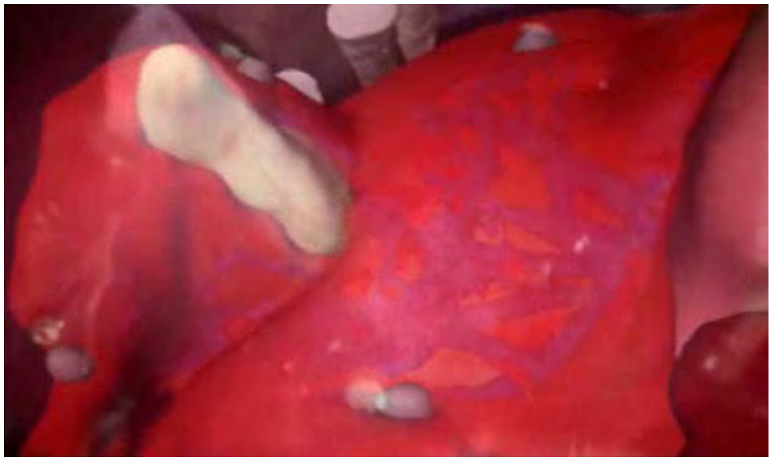

Figure 7 AR using 3D imagery from intraoperative CT. This novel AR technique preserves the surface texture information and optical depth cues from the laparoscopy while also exposing the underlying vasculature accurately.

4. DISCUSSION AND CONCLUSIONS

The described research is the first step toward an ambitious, long-term goal of taking advantage of volumetric imaging, now routine in diagnostic imaging, for minimally invasive surgeries. Laparoscopes, currently the primary visualization tool for navigating such surgeries, are limited in their 3D visualization capability. Essentially a video imaging technique, they cannot show structures under the exposed surfaces. AR, as proposed earlier, has provided the missing 3D information but is not accurate for extracranial surgeries, because the CT or MR imaging data employed for 3D visualization have not been current. We demonstrated the feasibility of using intraoperative data for AR here.

Our work combines real-time 3D imaging with minimally invasive laparoscopic surgery. Indeed, the greatest advantage of this approach is a dynamic 3D anatomical roadmap (rendering) to guide these procedures. Unlike the 3D roadmaps created in computer-assisted neurosurgery, our proposed 3D roadmap will refresh in real time to display tissue motion and surgical manipulations along with any surgical instruments within the operative field. Our approach is expected to initiate a new generation of minimally invasive surgeries relying on real-time 3D guidance. Incorporation of real-time 3D visualization and guidance is expected to allow laparoscopic surgeons to perform existing surgeries more precisely with fewer complications. Aided by improved visualization, it is also expected that many surgeries that are currently performed in an open invasive fashion can instead be performed minimally invasively, thereby reducing mortality and morbidity rates.

We have selected multislice CT as our intraoperative imaging modality for the proposed research. Any intraoperative imaging modality for guiding surgery must be volumetric and interactive (i.e., offer high frame rate). Current state-of-the-art 64-slice CT scanners provide a coverage of 4 cm, and we have been able to obtain 1 volumetric frame per second. Continuing advances in CT technology suggest that the coverage will grow to 10–12 cm. The potential exists for up to 8 volumetric frames per second speed as well. Therefore, evolving multislice CT technologies will become even more suitable for the proposed surgical application. In comparison, MR imaging lacks the speed needed for guiding interactive surgeries, and most surgical tools are not MR compatible. We have also found real-time 3D ultrasound unsuitable for the current application, because it cannot image across the pneumoperitoneum (caused by CO_2 insufflation). Cost and availability considerations also favor CT.

Radiation exposure to the patient and surgical team is a concern with the proposed continuous operation of the CT scanner. We showed up to 20-fold dose savings in an adult human through our strategy of elastically registering pre- and intraoperative CT images. Further savings may result if the CT scanners can be made to operate at even lower tube current setting. Visualization of critical underlying structures, especially the vasculature, is important before making surgical dissections. Inherent in our dose reduction strategy is a scheme to visualize the vessels without having to use the contrast continuously, which is neither permitted nor safe. Another advantage of having 3D models of anatomical structures in the CT-generated view is that one can interact with this view. For example, the surgeon can virtually practice a particular surgical manipulation and observe the effects of it in the CT view before actually making that manipulation. No such interaction is possible with the traditional laparoscopic view.

Engineering advances will be needed before continuous CT-guided laparoscopic surgery becomes routine. First, elastic image registration must be automatic and real-time. Our group has already made significant strides in this area, and we continue to improve the speed of elastic 3D image registration.[8,9] Full development of the proposed concept will also require a tight system-level integration among many subsystems and components: surgical tools, the laparoscope, tool tracking techniques, image processing (registration, in particular) technologies, visualization workstation, etc. Subsequently, a second level of integration will be required at the human/machine interface level that will combine the surgical protocol with the imaging protocol.

Continuous low-dose CT scanning of the dynamic operative field without exceeding acceptable radiation exposure and using high-speed elastic registration to generate diagnostic quality CT images of the intraoperative anatomy will enable high-quality 3D visualization of the operative field in a CT-equipped operating room. We have successfully created 3D renderings from multislice CT corresponding to anatomy visible within the field of view of the laparoscope. With additional developments, our research has the potential to help improve the precision of laparoscopic surgeries, reduce complications, and expand the scope of minimally invasive surgeries beyond its current 15% share of all surgeries.[2]

ACKNOWLEDGMENTS

This work was supported by the U.S. Department of Defense grant DAMD17-03-2-001.

REFERENCES

1. Himal HS. Minimally invasive (laparoscopic) surgery. *Surgical Endoscopy* 2002;16:1647-1652.
2. http://www.nibib1.nih.gov/events/IGI/IGIWorkshop2002_FINALReport.doc. 2002.
3. Harrell AG, Heniford BT. Minimally invasive abdominal surgery: lux et veritas past, present, and future. *Am J Surg* 2005;190:239-243.
4. Marescaux J, Rubino F, Arenas M, *et al.* Augmented-reality-assisted laparoscopic adrenalectomy. *Jama-Journal Of The American Medical Association* 2004;292:2214-2215.
5. Mutter D, Bouras G, Marescaux J. Digital technologies and quality improvement in cancer surgery. *Ejso* 2005;31:689-694.
6. Dandekar O, Walimbe V, Siddiqui K, Shekhar R, "Image registration accuracy with low-dose CT: How low can we go?", In Proceedings of 2006 IEEE ISBI; p 502-505.
7. Walimbe V, Shekhar R. Automatic elastic image registration by interpolation of 3D rotations and translations from discrete rigid-body transformations. *Medical Image Analysis* 2006; 10:899-914.
8. Castro-Pareja CR, Jagadeesh JM, Shekhar R., "FAIR: A hardware architecture for real-time 3-D image registration", *IEEE Trans Inf Technol Biomed* 2003;7(4):426-434.
9. Castro-Pareja CR, Shekhar R., "Hardware acceleration of mutual information-based 3D image registration", *J Imaging Sci Technol* 2005; 49(2):105-113.

A Pervasive Computing System for the Operating Room of the Future*

Sheetal Agarwal   Anupam Joshi   Tim Finin
Yelena Yesha
Department of Computer Science and Electrical Engineering
University of Maryland Baltimore County
{sheetal2, joshi, finin, yeyesha}@cs.umbc.edu

Abstract

We describe a prototype Context Aware Perioperative Information System to capture and interpret data in an operating room of the future. The captured data is used to construct the context of the surgical procedure and detect medically significant events. Such events, and other state information, are used to automatically construct an **E**lectronic **M**edical Encounter **R**ecord (EMR). The EMR records and correlates significant medical data and video streams with an inferred higher-level event model of the surgery. Information from sensors such as Radio Frequency Identification (RFID) tags provides basic context information including the presence of medical staff, devices, instruments and medication in the operating room (OR). Patient monitoring systems and sensors such as pulse oximeters and anesthesia machines provide continuous streams of physiological data. These low level data streams are processed to generate higher-level primitive events, such as a nurse entering the OR. A hierarchical knowledge-based event detection system correlates primitive events, patient data and workflow data to infer high-level events, such as the onset of anesthesia. The resulting EMR provides medical staff with a permanent record of the surgery that can be used for subsequent evaluation and training. The system can also be used to detect potentially significant errors. It seeks to automate some of the tasks done by nursing staff today that detracts from their ability to attend to the patient.

1 Introduction

Performing a surgery is an elaborate process and proceeds in progressive stages. The term Perioperative generally refers to the three phases of surgery namely preoperative, intraoperative, postoperative. The preoperative phase of a surgery involves identifying the patient, determining readiness of the nursing and medical staff, preparing the operating room, and capturing data from incoming medical records regarding vitals, pre-operative medications, tests, and scans etc. In a perioperative setting, hundreds of patients and staff may be flowing through dozens of operating rooms on a daily basis in a single facility. Some fraction of these patients are unscheduled and identified only on the day of surgery – this fraction can be large in facilities that deal with trauma. The resulting chaos can be overwhelming, even with some form of electronic health record (EHR) system (currently available in 12% of hospital systems). This is because the orchestration of behavior between information systems is people and paper based. Available automated systems are often dedicated to isolated operations or departments with no automated means to communicate with one other. The limitations of this environment provide great opportunity for process improvement efforts. It has been claimed by clinicians that a 30% improvement can be routinely achieved in almost any targeted area and 100% improvements are possible. At each site of care, the perioperative information

*This work supported in part by a subcontract from UM Medical School in context of the Traumapod Project

system must keep track of (and archive) what is done to each patient (both diagnostically and therapeutically), how, when, and why it is accomplished, and how the patient responds.

Our goal is to develop a *Context-Aware Perioperative Information System* that will automate most of the perioperative support functions such as patient tracking, inventory management, clinical documentation etc with the help of pervasive computing and semantic web technologies. In this paper we present the design of such a system, and a preliminary implementation that tracks supplies, prevents errors, and automatically creates an electronic medical encounter record to document the events occurring during a surgery.

We seek to create a system that can work in two environments. One is the “Operating Room of the Future”. This term is used to describe the concept that is sought to be realized by several surgical groups across the country and involves an IT enabled Operating Room. A good example is the ORF project at CIMIT (`http://www.cimit.org/orfuture.html`). The second is Trauma Pod [24], a DARPA project whose aim is to develop an automated surgical treatment system that does not require onsite medical personnel on the front lines of battle, and is ready to receive, assess, and stabilize wounded soldiers during the critical hours following injury.

The first phase of the program is an effort to develop robotic technology to perform a totally unmanned surgical procedure within a fixed facility. A human surgeon will conduct all the required surgical procedures from a remote location using a system of surgical manipulators. Automated robotic systems provide necessary support to the surgeon to conduct all phases of the operation. Our system seeks to provide the intelligence needed in the perioperative environment, that in a normal surgical setting is provided by the scrub and circulating nurses who are assisting the surgeon.

2 Background

Clinical record keeping in high velocity healthcare delivery environments like surgery is a necessary and critical task. It is also a time consuming task that detracts from hands on patient care and contributes to extraordinary labor costs associated with collecting, transcribing and re-keying records throughout the perioperative process. The details of the surgery are documented in patient charts called the *Perioperative Record* . This record contains information about the patients vital signs at periodic intervals, medicines administered, complications if any, supplies and tools used etc. Errors in medical documentation cost billions of dollars to the health industry every year [3]. Inaccurate records put not only the patient but also the healthcare provider at risk [25, 18]. Similarly, there is a need to track assets, both in terms of equipment and personnel. Perishable assets in particular need to be tracked as they are used, and the usage information needs to be integrated with supply chain management. There are studies that suggest significant costs are incurred in courriering medical supplies – to the tune of billions of dollar in a year. In addition, asset tracking can also help avoid surgical errors by ensuring that supplies such as sponges are properly dispensed and not left in the patient. Finally, there is a need to prevent fairly simple errors, such as the presence of incorrect equipment for the surgical procedure, or an incorrect surgeon or patient in the Operating Room.

The data recorded during the perioperative process become a part of the patients medical history and is used by physicians to give further treatment to the patient. Data collection in the operating room is complicated due to several reasons. Firstly, multiple providers (e.g.. surgeons, anesthesiologists, nurses) record data for a single care event (i.e., the patient’s surgery). Secondly, information collected by one provider is not readily available to another. Thirdly, experienced nurses assess the patients’ condition accurately and provide appropriate treatment, sometimes without documenting these procedures; thus, duplication or differences occur in documentation, data gathering can be cumbersome, and not all details are recorded. Moreover, in situations such as those envisaged in

Traumapod, there are no nurses or assistants – a remotely located surgeon is performing the procedure. Collecting data, maintaining the state of the system, and documenting the process therefore needs to be automated as well.

An Electronic Medical Encounter Record (EMR), has the potential to reduce documentation errors by minimizing data redundancy and providing accurate details of the ongoing surgery [2, 8]. Formally, an EMR is a medical record or any other information relating to the past, present or future physical and mental health, or condition of a patient, that resides in computers which process this data to deliver more efficient health-related services. The EMR is an essential part of systems like the Traumapod [24] where surgeries are performed by remotely controlled robots and no humans are involved in the process. Only the EMR can provide details of the events occurring during the surgery.

The operating room (OR) has several medical devices that provide information about the patients status. In addition to these devices, we can deploy sensors in the OR that can provide us with better view of the activities occurring in the operating room during a surgery. Unfortunately, most sensing technology detects low level events, and is rather error prone. For instance, consider the use of RFID tags on medicines. Clearly, the sensor will only tell us that the particular medicine is nearby or in the room – this information may or may not be valuable in of itself. What is more valuable is the knowledge that a particular medicine was administered at a given time. Moreover, in an OR with presence of fluids and metal, the accuracy of RFID sensing is itself not very high or robust.

We define a medically significant event as any event that affects or is a part of the surgical procedure. Many systems [27, 20, 28] have been built that monitor physiological parameters of a patient and signal alarming conditions. Healthcare providers use these alarms as cues as it is not possible to maintain a constant vigil over the patients' health status. The alarms are in the form of an audio alert or a message displayed on the computer screen that can be seen by the healthcare provider.

Most of these signal low-level events such as tachycardia, apnea or any other abnormal pathological state. As we will describe later, such low level events generally do not in of themselves provide any meaningful detail about the patients condition to the surgeon. To provide more meaningful information the alarms or medical events need to be interpreted at a higher level and documented. In addition to physiological data we can make use of data streams from sensors that can be deployed in an operating room to capture additional events such as tools and medicines used and identities of the members of the clinical staff. In our research we use the Radio Frequency Identification (RFID) system to detect medical supplies, tools and the staff.

3 Related Work

We developed a context-aware system that monitors and analyzes the data streams from various medical equipments and create an Electronic Medical Encounter Record, according to the inferences made by analyzing the data streams, in a perioperative environment. The surgical team can see the record being populated in real-time which ensures that everyone is aware of the progress being made and of the patients health status at all times. The system was designed to detect events during trauma care and general anesthesia scenarios.

Automated analyses of a patients' physiological data to detect alarming conditions has been a subject of research for over a decade. Several patient monitoring systems have been developed that alert the healthcare provider to alarming conditions. InCare [27],is one of the earliest automated systems to detect events in post-cardiac operated patients. InCare had a rule-based system that used multi-variable and trend based analysis of physiological data to detect events. Similarly, Schecke et al [28] designed a knowledge-based decision support system for patient monitoring in cardio anesthesia. The medications used and progress of the surgery was fed into the system manually by one of the members of the surgical staff.

Hewlett Packard Labs has recently developed a framework that allows development of scalable software systems to monitor and analyze continuous streams of data [5]. A prototype system BioStream was implemented to show its use in remote patient monitoring. BioStream is built on top of stream data processing architecture for real time processing of physiological signals. They use a database-oriented approach to analyzes data streams. The streams are subjected to "operators" that belong to a part of a patient plan. The current prototype is capable of identifying simple pathological conditions by monitoring ECG signals.

Bardram et al [10] developed a context-aware infrastructure to build context-aware applications for a hospital environment. The infrastructure includes sensors to detect presence of the nurse in the room, a smart pill container, a smart hospital bed to identify the patient. Radio Frequency Identification (RFID) is used to detect the people and the medications being used. A Context-Aware Electronic Patient Record was designed to present an user interface that adapts based on the current context. However, the work is in the preliminary stages and the focus is more on human computer interfaces.

Levine et al [29] have proposed the development of a computer automated perioperative situational awareness system that captures and records data from various medical devices and provides an integrated display to allow the operating team to visualize the data. The focus of this work so far has been on data capture and facilitating data visualization to provide context sensitive information and improve real time access to data. Our system focuses on making use of the data captured to detect medically significant events and creating an EMR.

Though the individual components of our system such as the algorithms to analyze physiological data, stream processing of data have been studied in previous systems, to the best of our knowledge no system has yet been developed to create an EMR in the perioperative environment.

4 System Architecture

The framework we developed is designed to collect, process and reason over surgical context information to detect medically significant events. It consists of a 3-tier event detection system. Events at the lower levels are processed to infer high-level events. At the lowest level, data is collected from various sensors in the operating room and processed by a stream processor to identify low level events like presence/absence of an RFID tag. These low level events are then processed by the second tier to detect low level medical events such as high blood pressure, apnea etc. Event detection at this level also makes use of patients medical history and knowledge of medicines used to detect events. At the highest level we have the rule based engine that correlates events from the second tier to detect medically significant events such as administration of anesthesia, tension pneumothorax etc.

4.1 Data Sources

- **Patient Monitors** The operating room has several patient monitoring systems that track the patients physiological parameters. For example, pulse oximeter monitors blood oxygen saturation levels, vitals signs monitors track heart rate, blood pressure etc. To monitor the patient's condition during the surgery, the surgical team monitors the value and change in physiological parameters. We use data streams from these patient monitors to determine the state of the patient during the surgery.
- **Asset and Personnel Tracking Systems** The system is designed to incorporate bluetooth and radio frequency identification to track medical supplies and nursing staff. The architecture can be easily extended to

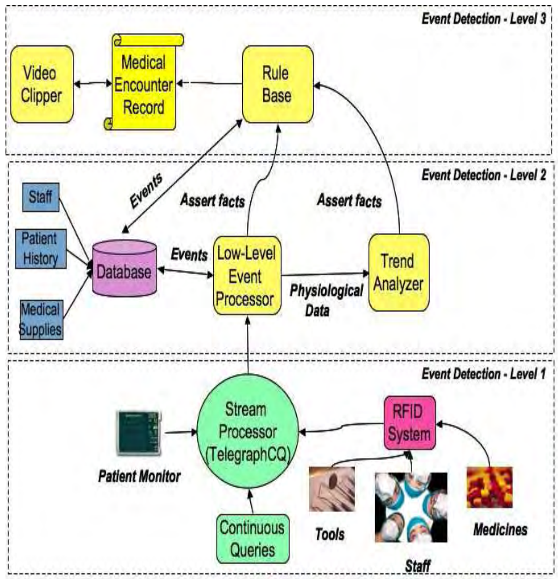

Figure 1: System Architecture

support other tracking technologies. For the prototype developed we used passive RFID tags. Readings from the RFID reader are analyzed to determine the resources used during the surgery and the team performing the surgery. Figure 2 shows the data sources used to acquire contextual information in the OR.

- **Message Exchanges** Where the operating room is highly automated, as in the Traumapod system, the nature and sequence of messages exchanged between its constituents can provide a rich source of state data. For example in Traumapod, we can listen to exchanges between the "Tool Rack" system and the scrub nurse system to figure out which tools have been moved from the rack to the (robotic) nurse. This can be done without any explicit sensing technology that would be needed in a live OR with human surgeons and nurses.

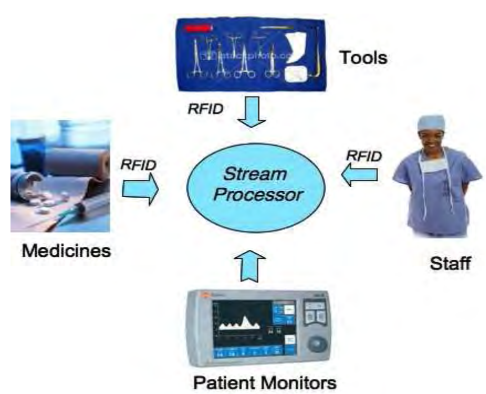

Figure 2: Data Sources in Operating Room

4.2 Low Level Event Management

The patient monitoring systems and the RFID reader produces continuous streams of data that need to be processed and analyzed in real-time to detect events. Traditional database systems have been designed to manage finite data sets where client queries are processed immediately against data stored in tables. In applications that processes continuous data streams, clients require long-running continuous queries that are evaluated as data streams through the application. For example, consider a query that monitors physiological data streams. ”Report the heart rate and blood pressure values when heart rate is ¿ 100 an blood pressure ¡ 70 over a period of 60 seconds”. Significant work has been done is the area of stream processors. Some of the well know systems are Stanford STREAM data manager,Aurora, Borealis and TelegraphCQ. Each of these systems focus on processing continuous time-varying data.

In our prototype implementation we used, TelegraphCQ [26], developed at University of California, Berkeley to process the physiological and RFID data streams. Data from patient monitors and the RFID reader is pushed to the stream engine continuously. Queries over these data streams are specified over a time window. As new data arrives, the queries are evaluated and results are returned to the client.

4.3 Analyzing RFID Data streams

Figure 3 shows the RFID system we used in our prototype. We used the Symbol AR400 900MHz Reader and passive RFID tags. An RFID tag has a unique 96-bit identifier called the Electronic Product Code (EPC). The RFID reader returns the list of EPC codes it detects. We implemented the *Byte Stream Protocol* to interface with the RFID reader. The RFID API we developed provides a layer of abstraction over the low level protocol. The API processes results from the RFID reader.

Figure 3: RFID System

The RFID module polls the reader periodically to get the list of RFID tags visible. The passive tags use the energy incident from the reader, to return their EPC code. The reader reads the tags at its own internal frequency. Hence the same tag may be reported more than once in the list of tags detected. When many tags are in close proximity, the signals returned by the tag collide and result in loss of data. Thus a single read from the reader is not sufficient to detect all tags reliably. A continuous query over the RFID data stream, aggregates the number of times a particular tag is reported by the reader.The reader is sampled every 2 seconds.Experiments show that if a tag is detect at least 5 times in a 30 seconds window i.e. 66.67% of the time, then the tag is visible.

```
If Number of Times Tag Seen >= 5 then
   Event (Tag Visible)


If tag not seen for >= 120 seconds then
   Event (Tag Invisible)
```

4.4 Analyzing Traumapod Messages

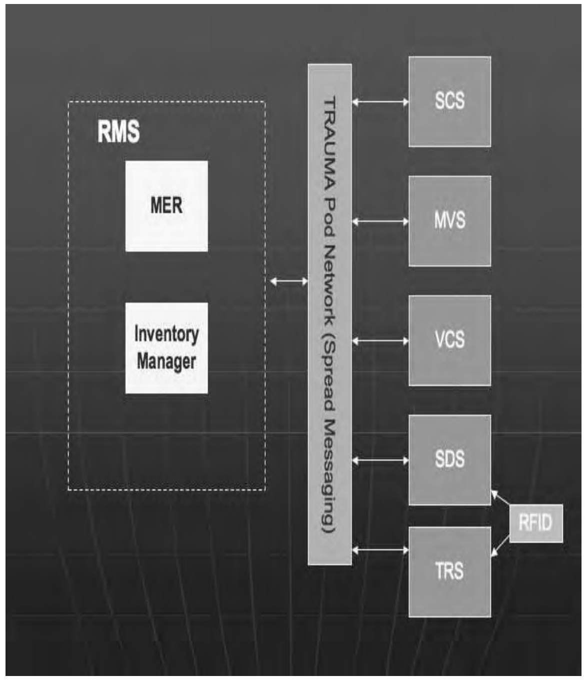

Figure 4: Electronic Medical Encounter Record

As mentioned earlier, the sequence of message between components of Traumapod can also be used to infer state and create context. Figure 4 shows the traumapod network. The various systems in the network communicate using the Spread messaging toolkit[1]. This is an open source messaging system developed at Johns Hopkins University. Each system in the network implements a set of interfaces. Some interfaces are implemented by all subsystems and each subsystem has its own interface defined.

The Supply Dispensing System(SDS) is a robotic system that contains the medical supplies in form of supply trays. This system is responsible for dispensing the supplies requested by the surgeon. The Tool Rack System (TRS) is a robotic system that contains the surgical tools. This system dispenses the tools required during the surgery. The Scrub Nurse System (SNS) moves the supplies and tools between the surgical site and the SDS or TRS. The Supervisory Control Subsystem (SCS) is the master system that controls all the other subsystems. Any request for a tool or a supply is first sent to the SCS. The SCS then executes the task by sending appropriate commands to SDS, SNS and TRS.

The component we add to Traumapod is called the Resource Management System (RMS). It is a passive system that snoops for messages between SCS, SDS and TRS and based on the messages infers the state of the surgery. For each procedure that is to be performed , the surgeons typically define a surgical template outlining the key steps in the surgery. These steps are broken down into a sequence of messages that would be expected if that step was being performed. For instance, a surgeon will give a voice command to the system request some sutures. This would be directed to the SCS, which would then send messages to the SDS to dispense that suture, and to the SNS to transfer the suture to the surgeon. The messaging template is used as only a hint. The actual surgery may involve additional steps. A rule base is created to map message sequences to a surgical step, and the steps to key events of the surgery. The rules are based on traumapod events such as supply dispensing, tool requests and

[1] `http://www.spread.org/`

tool changes. The current rule base consists of rules to detect events during a “Shunt Procedure”, which is the procedure that the initial phase of the Traumapod project is expected to demonstrate.

5 Detecting Medically Significant Events

The highest level of event detection in the system consists of a rule-based system, JESS [11] with FuzzyJ library. This system define the rules to detect the medically significant events. The knowledge base was developed by gathering information from an anesthesiologist by interviews and from medical literature [17, 14, 13, 12, 19, 15, 16] which describe methods to analyze and interpret physiological data.

Physiological parameters reflect a patients’ health status. Interpretation of physiological data to infer the patients condition is a challenging problem. Some of the earliest data analysis systems used simple limits on physiological parameters for basic interpretation. The more advanced systems considered dependencies between parameters to provide more meaningful interpretations [27, 20, 28] . The problem with such systems is the high rate of false positive or false negative events. The poor performance was due to the fact that physiological data was interpreted independently of the clinical conditions in which the data was acquired. Physiological parameters not only depend on the physiologic processes but also on factors such as patients current condition, medical history, medicines administered and the sequence of occurrence of other events. Most systems consider only a subset of these factors. Since our system can directly detect or infer many of these factors, we can do a better job of interpreting the parameters.

Each physiological parameter has a range of values that can be classified as normal or abnormal. However, given a data value, there is no set threshold that will deterministically classify the value as normal or abnormal. Also as mentioned above the interpretation of a parameter varies with the clinical context. We use the fuzzy set theory to capture this uncertainty in medical data. Fuzzy membership functions are used to classify data values. The value can be ”very low”, ”low”, ”normal”, ”high” or ”very high”. Rate of change is another important factor that is used to determine the health status of the patient. The change can be ”constant”, ”stable” or ”abrupt” and the value can be ”increasing” or ”decreasing”. The value and rate of change of value is used to detect events.

The value of the membership varies between 0 and 1 where 1 implies absolute membership. The set points used to define the range of values varies with each patient. For example, the range of normal blood pressure for a hypotensive patient will be different than the range for a patient with normal blood pressure. In the current version of the system, the set point for each parameter is preset for a patient. In ongoing work we are extending the system to set these limits by analyzing the patients medical history and the pre-op diagnosis.

The membership functions used for each parameter were different and partitioning of the range of values was determined by eliciting information through interviews with an anesthesiologist. Some of the functions used were [9] TriangleFuzzySet, TrapezoidFuzzySet, SFuzzySet etc. To defuzzify the values we use the maximum defuzzification function. In this method the mean of the x values, with maximum membership values over the entire set of FuzzyValues, is calculated.

Given the data value of a physiological parameter and its trend, the techniques we used to correlate low level events are:

- **Pre-op Diagnosis**

  Typically, before the patient is brought into the surgery, the patient’s condition is evaluated. The evaluation includes taking note of the vital signs, any medical care provided, a physical examination and any other notable medical condition. The pre-op diagnosis is used to initialize the event history. The actions taken during the surgery also depend on the pre-op diagnosis, so we used this information in detecting events. For example, if the patient was bleeding excessively prior to the surgery, detecting fluid infusions during the surgery is more accurate. Currently we use pre-op diagnosis to only provide clues about the patients' condition before the surgery starts.

- **Multi-variable Analysis**

  Monitoring a physiological parameter in isolation does not give much information about the state of the patient. Coleman et al. [30] state that "each physiologic state variable is intimately related directly and indirectly to many others by relationships that depend on the condition of the subject". This means that physiologic parameters not only depend on physiologic processes but are also affected by the patients current condition. For example low and decreasing blood pressure, does not give signify too much detail. However, low and decreasing blood pressure with high and increasing heart rate implies potential loss of fluids. Monitoring a physiological parameter along with its relationship with other parameters helps determine more meaningful events.

- **Event History** Event history is a set of low-level and high-level events already detected. A high level event is often a composition of low-level events and potentially other high level events. The composition can be a conjunction or disjunction of events [6]. Often, events need to be considered in the context of the event history and both a high and low level.

  Example: Event Conjunction

  ```
  If (TensionPneumothorax)
     and (Systolic BP "low" and
     "increasing")
     and (Heart Rate "high"
     and "decreasing") then
         Event (Decompression)
  ```

- **Effect of Medicines**

  The medicines administered during the surgery, may or may not have a significant effect on the patient's physiology. The time the medicine was administered, effects expected, duration of effect and time to affect are some of the factors that need to be taken into account to detect their effect in the physiological parameters. Determining the event of medicine administration is extremely difficult without the context of the medicine used. In our system, the medical supplies are tagged using RFID tags. Thus we can acquire the contextual information using which we can detect medicine administration. Detection of the medicine does not imply if the medicine was actually administered. In the current version of the system we record all the medicines detected by the RFID system. For those medicines whose effect is observed in the physiological parameters, an event is signaled to indicate that the medicine was actually used. Incorporating the information of medications used to detect events is a difficult problem. Firstly, strength and duration of the effect of the medication may vary with each individual. Secondly, it is difficult to estimate the adverse effect a medicine may have on a patient. We have preliminary results to show the utility of such information for event detection. In some cases, it has been suggested that a dedicated short range RFID type sensor exist and that the medicine to be administered be flashed before it to indicate administration. Where practicable, such a system can be of significant help.

  Example:

```
If (Respiratory Rate "abrupt decrease"
   and Systolic BP "abrupt decrease"
   and Heart Rate "stable")
       If (Time Anesthetic detected < 40 sec)
          then Event (Start Anesthesia)
```

6 Results

We have created a prototype of the system described above and evaluated it in a surgical training setting. We first describe the EMR Interface that is actually available to the surgeons and trainees to interact with our system, and then show results about how accurate our event detection system is.

6.1 System Interface: Electronic Medical Record

The EMR can be displayed on a computer screen in the operating room. It provides a summary of the patient profile, the pre-op diagnosis, laboratory reports, and radiology/imaging reports. Before surgery can commence, obvious errors such as right patient, right surgeon, right equipment and the presence of consent form (assumed to be RFID tagged) are checked. Any noted errors are announced. If no errors are detected, permission to proceed is provided.

During surgery, the vital signs of the patient are streamed to the screen during the surgery. The event list gets populated as events are detected. Typically, only clinically significant events are populated on the screen. A part of the screen is used to show the medicines and the surgical staff as detected by the RFID system. As members of the surgical team enter and leave the OR, the screen is updated to show only those present in the room. For each event we save the vital signs of the patient at that instant of time. In many surgical situations videotaping is done. Our system can accept and display the feed. In addition, post surgery, video clips of defined length for each of the clinically significant events are created and the corresponding video URL is stored in the medical encounter record. When reviewed at any time after the surgery, the surgeon or a trainee can interact with the record to see the vital signs of the patient at the major points in the surgery and the associated video data. Instead of viewing the entire video footage of the surgery, the surgeon can browse through the key parts of the video by selecting an event from the event list. Figure 5 shows a snap shot of the EMR.

In addition, separate interfaces are available that allow access to the state information relating to supplies that the system maintains. In other words we can see what supplies and tools were used during the procedure, and what numbers remain. We hope that this information can be subsequently integrated into the supply chain management system at the hospital.

6.2 Event Detection Accuracy

In this section we describe the test environment we used to evaluate our system. We used physiological data sets from the Human Patient Simulator (HPS)[2] called Stan, manufactured by the METI Inc. It is a complex system that emulates the human body response to medical treatment. The simulator is used to train medical students.

[2] http://www.meti.com/

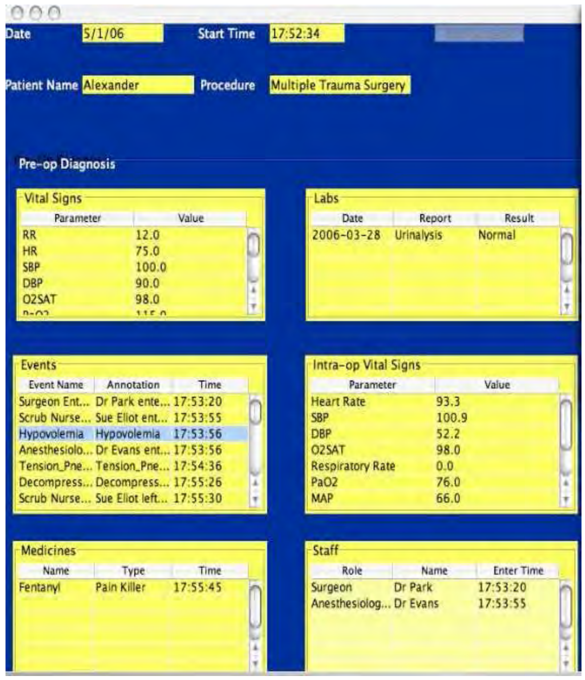

Figure 5: Electronic Medical Encounter Record

Stan, shown in Figure 6, can be loaded with various patient profiles. For example, the doctor could create an asthmatic patient with chronic heart disease who is taking a handful of certain drugs and is currently experiencing anaphylactic shock, a severe allergic reaction. The medical students in turn have to figure out how to treat the patient. If medication is required, the drugs are "administered" by scanning a bar code on a syringe. The computer produces in Stan the physiological response that the drug would have produced in a patient with that medical condition.

This system was made available to us by the Air Force Simulation Center at University of Maryland Medical School. In order to evaluate our system we used two custom scenarios. Since Traumapod [24] focuses on trauma care on the battlefield, we chose to use trauma related scenarios to evaluate our system. The HPS remains in each of the states in a given scenario for a fixed period of time after which it transitions to the next state. The changes in the physiological parameters of the simulator are logged constantly and the parameters vary according to the current state of the HPS. Several variants of the trauma scenario were tested, we describe one particular scenario here.

Figure 6: Human Patient Simulator

Scenario : Blunt Trauma Multiple Injuries

This scenario consists of a patient who has been wounded in a battlefield. In this scenario the patient is goes through the following states during the course of trauma care:

- Hypovolemia (Excess blood loss)
- Tension pneumothorax
- Decompression
- Fluid Infusions

These states constitute the set of medically significant events we wish to detect from low level sensor data stream. This scenario was simulated on different patient profiles. Each patient has different medical history and pre-op diagnosis. Each of the scenarios was simulated with five different profiles. Slight variations of the scenarios were simulated to give us more varied data sets. The scenarios chosen each had the possibility of false positives and negatives. For instance, conditions to detect hypovolemia are increasing heart rate and a decreasing blood pressure. However, these conditions occur during Tension pneumothorax also. One key criterion of success is low false positives (and negatives). We found that for most events, we had very high true positives (in the 90+ percent range), and very low false negatives. The exception was events such as Hypovolemia and Fluid Infusion, which depend on conditions for which we do not have all vitals available. However, even here the use of event history and pre-op diagnosis significantly increased the true positives over just using the available vitals.

The other concern is the latency between the occurrence of an event and its detection by the monitoring algorithm plays an important role in the performance of the system, as does the order of detection. We detected all events in proper order. The latency varied from a few seconds all the way upto 56 seconds, and the variance was generally high. This was especially true for events where the response of the vitals "develops" over time and is not instantaneous.

7 Discussion

Our rule base currently has 27 rules. Adding and retracting facts from the knowledge base is an expensive operation. We designed the knowledge base to minimize such operations. As rfid events are detected facts are either asserted or retracted. For our rule base we start with a knowledge base of 12 initial facts.

The system evaluation is extremely positive, with many events having 100% true positive and almost zero false positives . However, admittedly the data set was small and was obtained from a Human Patient Simulator. The results may vary with real patient data. They will also likely vary as more scenarios are evaluated. Also as the knowledge-base grows, addition of new rules to detect more events may increase the number of false positives.

Currently we use simple queries over the data streams to detect low-level events. We maintain a state variable model and use various techniques to correlate these low-level events to infer more meaningful events. Some of the event correlation can be done by using appropriate queries on data streams. The current version of TelegraphCQ does not provide support for sub-queries and access to historical data. With support for sub-queries and access to archived data in the subsequent version we can move some of the event correlation rules to the stream processing level.

We use RFID to detect staff and medicines in the operating room. The use of RFID in healthcare presents a number of critical issue unique healthcare in addition to the basic limitations of the technology.

- **Electromagnetic Interference:** The healthcare environment is already full of safety critical devices that are sensitive to radiation at various frequencies.
- **Tagging Medical Supplies:** We conducted a feasibility study of using RFID to tag medical supplies. The current state of art is not sophisticated enough to allow tagging of all medical supplies. The smallest passive tags available are 1" x 1". With tags of this size it is difficult to tag items like surgical tools, medical supplies like cotton balls, sponges, gauze etc. Tags that are of the size of a grain of rice are also available. But these tags are designed to embed under the skin of cattle or humans. These are not suitable to tag medical supplies.
- **Environment Hazards to Tags:** The healthcare industry presents a unique challenge to the physical integrity of RFID tags because of its pervasive infection control measures. Supplies like sponges, gauze become wet with fluids. Tags attached to clothes may be damaged when they are washed. The RFID tags were originally designed to tag objects for supply chain management and are not capable of withstanding harsh medical environments.

Inspite of the above the limitations, the use of RFID in healthcare is expected to rise rapidly. According to a report published by IDTechEx [22] "the market for RFID tags and systems in healthcare will rise from $90 million in 2006 to $2.1 billion in 2016. Primarily, this will be because of item level tagging of drugs and Real Time Locating Systems (RTLS) for staff, patients and assets to improve efficiency, safety and availability and to reduce losses." The technology is expected to evolve to address the requirements of RFID in hospital environments.

Another important aspect to consider is that the EMR documents only those events that are detected by the system. A complete perioperative record has several details such as physical observations of the patients body, devices implanted, exact amounts of fluids infused etc. These are the kind of details that cannot be deduced from the data sources we currently use.

8 Ongoing Work

In this section we describe some of the ongoing extensions we are implementing and enhancements that could improve the system.

8.1 Domain-Based Medical Ontology

A knowledge-based system represents relationships between objects, entities and concepts that exist in a domain of interest. Ontology is a specification of such concepts. The relationship between the objects is specified in a vocabulary that is used by the knowledge systems to represent knowledge [23]. Within health informatics, ontology is a formal description of a health-related domain.

The use of ontologies in medicine is mainly focused on the representation and (re-)organization of medical terminologies. Physicians developed their own specialized languages and lexicons to help them store and communicate general medical knowledge and patient-related information efficiently. Such terminologies, optimized for human processing, are characterized by a significant amount of implicit knowledge. Medical information systems, on the other hand, need to be able to communicate complex and detailed medical concepts (possibly expressed in different languages) unambiguously.

In the perioperative environment, use of a standardized language decreases patients' risk for injury by eliminating inconsistency of language or meaning. This is a difficult task and requires a detailed analysis of the structure and the concepts of medical terminologies. But it can be achieved by constructing medical domain ontologies for representing medical terminology systems.

The benefits of using a medical ontology are:

- Ontologies can help build more powerful and more interoperable information systems in healthcare.
- Ontologies can support the need of the healthcare process to transmit, re-use and share patient data.

Constructing the medical encounter record using a domain-based ontology will make the record usable by other health-informatics systems for further processing. Several groups, such as GALEN [7], CIMIT [1], SNOWMED-CT [4], have developed medical ontologies to represent medical concepts. Most groups focus on a domain within medicine and have their ontology represent concepts relevant to the domain. The Unified Medical Language System (UMLS) [21] is a meta-thesaurus created by the National Library of Medicine (NLM) that integrates the ontologies developed by various groups.

Moreover, ontologies and Semantic Web languages such as OWL can be used to capture domain knowledge and rules explicitly and in a machine interpretable way. This will allow for significantly greater interoperability at a semantic level.

8.2 Fine grained Tracking with RFID

Supply counting is an important procedure during a surgery. It is the responsibility of the surgical team to ensure that no supply is left within the patients body at the end of the surgery. RFID can be used to perform supply counts provided all supplies can be tagged. Since RFID tags cannot be localized, as an alternate solution we can use low frequency readers to detect tags in a particular zone of the operating room. The ability to divide the operating room in zones will allow us to track the supplies in the operating room and ensure no supply is left within the patient's body.

Tracking supplies at this granularity can also be useful in inferring events that are not detectable through physiological data streams. For example, if the surgeon is holding a vascular clamp and the surgery involves placing a shunt, we can estimate the time that the clamp was used to clamp the blood vessels. With the current system, such events are not detectable.

9 Conclusion

We presented a prototype of a context-aware system that analyzes data streams in an operating room to detect medically significant events and document them in an electronic encounter record. The system uses technologies like Radio Frequency Identification to acquire contextual information such as resources used and the staff present in the OR. We explored the use of medical history and effect of medicines on physiology and we conclude that such techniques help us detect complex and more meaningful medical events. The system architecture is scalable and can be extended to detect events over larger number of scenarios.

10 Acknowledgements

We would like to thank Colonel Grissom at the Air Force Simulation Center, University of Maryland Medical School, for the use of the HPS METI system to evaluate our system. We would also like to thank Dr. Geoffrey Kuzmich for the invaluable clinical advice on developing the medical rule base without which this work would not be possible. We would also like to thank Dr Franklin and his students at University of California Berkeley, for their support and expertise in using TelegraphCQ.

References

[1] CIMIT, Center For integration of MEDICINE and Innovative Technology. http://www.cimit.org/, Accessed April 2006.

[2] Electronic Medical Records. http://www.expert-system.com/medical_errors.htm, Accessed April 2006.

[3] Medical Errors: The Scope of the Problem. Technical report, Agency for HealthCare Research and Quality. Available at: http://www.ahrq.gov/qual/errback.htm.

[4] SNOMED-CT. http://snomed.org/, Accessed April 2006.

[5] Bar-Or A., Goddeau D., Healey J., Kontothanassis L., Logan B., Nelson A., and Thong J. BioStream: A System Architecture for Real-Time Processing of Physiological Signals. Technical report, Hewlett Packard Labs, 2004.

[6] Nagargadde A., Srihar V., and Ramamritham K. Representation and Processing of Information Related to Real World Events. Applied Research Group, Satyam Computer Services Ltd.

[7] Rector A. and Rogers J. The GALEN ontology. In *Medical Informatics Europe (MIE 96)*, 1996.

[8] Briggs B. Electronic Records: Protection on the Road to Patient Safety. *Health Data Management*, 2003.

[9] Orchard B. Fuzzy Sets. http://www.iit.nrc.ca/IR_public/fuzzy/fuzzyJDocs/FuzzySet.html, Accessed April 2006.

[10] Bardram J. E. Applications of Context-Aware Computing in Hospital Work Examples and Design Principles. In *Proceedings of the 2004 ACM Symposium on Applied Computing*, 2004.

[11] Friedman-Hill E. Jess, the Rule Engine for the Java Platform. *Distributed Computing Systems*, 2003.

[12] Hunter J. and McIntosh N. Knowledge-Based Event Detection in Complex Time Series Data. In *AIMDM '99: Proceedings of the Joint European Conference on Artificial Intelligence in Medicine and Medical Decision Making*, pages 271–280, London, UK, 1999. Springer-Verlag.

[13] Roberts J., Stalow C., and Hedges J. *Clinical Procedures in Emergency Medicine*. Elsevier, 2005.

[14] Sukuvaara T. Koski E., MŁkivirta A. and Kari A. Clinicians' Opinions on Alarm Limits and Urgency of Therapeutic Responses. *Journal of Clinical Monitoring and Computing*, 12(2):85–88, 1996.

[15] Tsien C. L. Event Discovery in Medical Time Series Data. In *American Medical Informatics Association (AMIA) Symposium*, 2000.

[16] Tsien C. L. Trend Finder: Automated Detection of Alarmable Trends. Technical report, Massachusetts Institute of Technology., 2000.

[17] Bloom M. Techniques to Identify Clinical Contexts During Automated Data Analysis. *Journal of Clinical Monitoring and Computing*, 10(1):17–22, 1994.

[18] Edwards M. and Moczygemba J. Reducing Medical Errors Through Better Documentation. *The Health Care Manager*, pages 329–333, 2004.

[19] Krol M. and Reich D. L. The Algorithm for Detecting Critical Conditions During Anesthesia. In *CBMS '99: Proceedings of the 12th IEEE Symposium on Computer-Based Medical Systems*, page 208, Washington, DC, USA, 1999. IEEE Computer Society.

[20] Navabi M., Mylrea K., and Watt R. Detection Of False Alarms Using An Integrated Anesthesia monitor. In *IEEE Engineering in Medicine and Biology Society 11th Annual Conference*, 1989.

[21] Bodenreider O. The Unified Medical Language System (UMLS): integrating biomedical terminology. *Nucleic Acids Research*, 32:267–270, 2004.

[22] Harrop P. and Das R. RFID in Healthcare 2006-2016. Technical report, IDTechEx.

[23] Corazzon R. Ontology, A Resource Guide for Philosophers. http://www.formalontology.it, Accessed March 2006.

[24] Satava R. Telesurgery, robotics, and the future of telemedicine. *European Surgery*, 37(5):304–307, 2005.

[25] J. Roediger and P. Salmon. Making Changes in Charts: Do's and Don'ts. `http://www.muschealth.com/professionals/ppd/chartshanges.htm`, Accessed November 2006.

[26] Chandrasekaran S. TelegraphCQ: Continuous Dataflow Processing for an Uncertain World, 2003.

[27] Mkivirta A. Sukuvaara T., Koski E. and Kari A. A Knowledge-Based Alarm System for Monitoring Cardiac Operated Patients : Technical Construction and Evaluation. *Journal of Clinical Monitoring and Computing*, 10(2):117–126, 1993.

[28] Schecke T., Langen T., Popp H., Rau G., KŁsmacher H., and Kalff G. Knowledge-Based Decision Support for Patient Monitoring in Cardioanesthesia. *Journal of Clinical Monitoring and Computing*, 10(1):1–11, 1993.

[29] Levine W., Meyer M., Brzezinski P., Robbins J., and Sandberg W. Computer Automated Total Perioperative Situational Awareness and Safety Systems. *Computer Assisted Radiology and Surgery*, 2005.

[30] Coleman WP., Seigel JH., Giovanni I., DeGaetano A., Goordarzi S., and Tacchino RM. Probability and Patient State Space. *Journal of Clinical Monitoring and Computing*, pages 201–215, 1990.